



---

## ChorusOS man pages section 2SEG: Virtual Memory Segment Services

---

Sun Microsystems, Inc.  
901 San Antonio Road  
Palo Alto, CA 94303-4900  
U.S.A.

Part No: 806-3330  
December 10, 1999

Copyright 1999 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, ChorusOS, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

**RESTRICTED RIGHTS:** Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 1999 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, Californie 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, ChorusOS, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



# Contents

---

## **PREFACE 7**

dcAlloc(2SEG) 13

dcFree(2SEG) 13

dcCluster(2SEG) 15

dcFillZero(2SEG) 16

dcSync(2SEG) 17

dcFlush(2SEG) 17

dcAlloc(2SEG) 20

dcFree(2SEG) 20

dcGetPages(2SEG) 22

dcIsDirty(2SEG) 24

dcPgNumber(2SEG) 25

dcPxmDeclare(2SEG) 26

dcRead(2SEG) 27

dcWrite(2SEG) 27

dcSync(2SEG) 29

dcFlush(2SEG) 29

dcTrunc(2SEG) 32

dcRead(2SEG) 33

dcWrite(2SEG)	33
lcOpen(2SEG)	35
lcClose(2SEG)	35
lcCap(2SEG)	35
lcOpen(2SEG)	36
lcClose(2SEG)	36
lcCap(2SEG)	36
lcFillZero(2SEG)	37
lcFlush(2SEG)	39
lcSetRights(2SEG)	39
sgFlush(2SEG)	39
sgSyncAll(2SEG)	39
vmFlush(2SEG)	39
lcOpen(2SEG)	42
lcClose(2SEG)	42
lcCap(2SEG)	42
lcPushData(2SEG)	43
lcRead(2SEG)	45
lcWrite(2SEG)	45
lcFlush(2SEG)	47
lcSetRights(2SEG)	47
sgFlush(2SEG)	47
sgSyncAll(2SEG)	47
vmFlush(2SEG)	47
lcStat(2SEG)	50
sgStat(2SEG)	50
lcTrunc(2SEG)	52
lcRead(2SEG)	54

lcWrite(2SEG)	54
MpCreate(2SEG)	56
MpGetAccess(2SEG)	57
MpPullIn(2SEG)	61
MpPushOut(2SEG)	64
MpRelease(2SEG)	66
pageIoDone(2SEG)	67
pageMap(2SEG)	68
pageUnmap(2SEG)	68
pagePhysAddr(2SEG)	70
pageSetDirty(2SEG)	71
pageSgId(2SEG)	72
pageMap(2SEG)	73
pageUnmap(2SEG)	73
PxmOpen(2SEG)	75
PxmClose(2SEG)	75
PxmGetAcc(2SEG)	77
PxmOpen(2SEG)	80
PxmClose(2SEG)	80
PxmPullIn(2SEG)	82
PxmPushOutAsyn(2SEG)	83
PxmRelAccLock(2SEG)	85
PxmStat(2SEG)	86
PxmSwapOut(2SEG)	87
rgnFlush(2SEG)	89
rgnInit(2SEG)	91
rgnInitFromDtCache(2SEG)	93
rgnMap(2SEG)	95

rgnMapFromDtCache(2SEG) 97

lcFlush(2SEG) 99

lcSetRights(2SEG) 99

sgFlush(2SEG) 99

sgSyncAll(2SEG) 99

vmFlush(2SEG) 99

sgRead(2SEG) 102

sgWrite(2SEG) 102

lcStat(2SEG) 104

sgStat(2SEG) 104

lcFlush(2SEG) 106

lcSetRights(2SEG) 106

sgFlush(2SEG) 106

sgSyncAll(2SEG) 106

vmFlush(2SEG) 106

sgRead(2SEG) 109

sgWrite(2SEG) 109

lcFlush(2SEG) 111

lcSetRights(2SEG) 111

sgFlush(2SEG) 111

sgSyncAll(2SEG) 111

vmFlush(2SEG) 111

**Index 113**

# PREFACE

---

---

## Overview

A man page is provided for both the naive user, and sophisticated user who is familiar with the ChorusOS™ operating system and is in need of on-line information. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

The following is a list of sections in the ChorusOS man pages and the information it references:

- *Section 1CC: User Utilities; Host and Target Utilities*
- *Section 1M: System Management Utilities*
- *Section 2DL: System Calls; Data Link Services*
- *Section 2K: System Calls; Kernel Services*
- *Section 2MON: System Calls; Monitoring Services*
- *Section 2POSIX: System Calls; POSIX System Calls*
- *Section 2RESTART: System Calls; Hot Restart and Persistent Memory*
- *Section 2SEG: System Calls; Virtual Memory Segment Services*
- *Section 3FTPD: Libraries; FTP Daemon*
- *Section 3M: Libraries; Mathematical Libraries*
- *Section 3POSIX: Libraries; POSIX Library Functions*
- *Section 3RPC: Libraries; RPC Services*
- *Section 3STDC: Libraries; Standard C Library Functions*
- *Section 3TELD: Libraries; Telnet Services*
- *Section 4CC: Files*

- *Section 5FEA: ChorusOS Features and APIs*
- *Section 7P: Protocols*
- *Section 7S: Services*
- *Section 9DDI: Device Driver Interfaces*
- *Section 9DKI: Driver to Kernel Interface*
- *Section 9DRV: Driver Implementations*

ChorusOS man pages are grouped in Reference Manuals, with one reference manual per section.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the `intro` pages for more information and detail about each section, and `man(1)` for more information about man pages in general.

NAME	This section gives the names of the commands or functions documented, followed by a brief description of what they do.
SYNOPSIS	<p>This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full pathname is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.</p> <p>The following special characters are used in this section:</p> <ul style="list-style-type: none"> <li>[ ]     The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.</li> <li>. . .    Ellipses. Several values may be provided for the previous argument, or the previous argument can be specified multiple times, for example, 'filename . . .'.</li> <li>         Separator. Only one of the arguments separated by this character can be specified at time.</li> <li>{ }     Braces. The options and/or arguments enclosed within braces are</li> </ul>



interdependent, such that everything enclosed must be treated as a unit.

FEATURES	This section provides the list of features which offer an interface. An API may be associated with one or more system features. The interface will be available if one of the associated features has been configured.
DESCRIPTION	This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES.. Interactive commands, subcommands, requests, macros, functions and such, are described under USAGE.
OPTIONS	This lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.
OPERANDS	This section lists the command operands and describes how they affect the actions of the command.
OUTPUT	This section describes the output - standard output, standard error, or output files - generated by the command.
RETURN VALUES	If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.
ERRORS	On failure, most functions place an error code in the global variable <code>errno</code> indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.

USAGE	This section is provided as a guidance on use. This section lists special rules, features and commands that require in-depth explanations. The subsections listed below are used to explain built-in functionality:
EXAMPLES	<p>Commands Modifiers Variables Expressions Input Grammar</p> <p>This section provides examples of usage or of how to use a command or function. Wherever possible, a complete example including command line entry and machine response is shown. Whenever an example is given, the prompt is shown as <code>example%</code> or if the user must be superuser, <code>example#</code>. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS and USAGE sections.</p>
ENVIRONMENT VARIABLES	This section lists any environment variables that the command or function affects, followed by a brief description of the effect.
EXIT STATUS	This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion and values other than zero for various error conditions.
FILES	This section lists all filenames referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.
SEE ALSO	This section lists references to other man pages, in-house documentation and outside publications.
DIAGNOSTICS	This section lists diagnostic messages with a brief explanation of the condition causing the error.
WARNINGS	This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.
NOTES	This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.

## BUGS

This section describes known bugs and wherever possible, suggests workarounds.

# Virtual Memory Segment Services

<b>NAME</b>	dcAlloc, dcFree – Allocate a data cache for a segment; Free a previously allocated data cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcAlloc(KnSgId sgId, KnExtPxMapper * pxm, VmFlags flags, KnLcId * lcIdp);  void dcFree(KnLcId lcId);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcAlloc</i> function allocates a new data cache object associated with the data segment identified by the <i>sgId</i> parameter. Two subsequent calls to <i>dcAlloc</i> with the same segment identifier will allocate two different data cache objects. It is thus the responsibility of the invoker to insure a one-to-one mapping between data cache objects and data segments. The <i>sgId</i> is an opaque value for the nucleus which is only used as an argument within upcalls from the VM to the external Proxy-Mapper managing the data segment. The <i>pxm</i> argument specifies the routines which will be used to perform upcalls required for this data segment (see <i>dcPxmDeclare(2SEG)</i>). The <i>flags</i> argument specifies the required properties of the data segment. It must be either <code>K_NOSWAPOUT</code> if the data cache must not be swapped, or 0. If successful, <i>dcAlloc</i> returns <code>K_OK</code> and sets the <i>lcIdp</i> output argument to the identifier of the newly allocated data cache object, otherwise an error code is returned and the value of the <i>lcIdp</i> field is undefined. The <i>lcIdp</i> returned by <i>dcAlloc</i> is an opaque for the External Proxy-Mapper and must only be used as an argument for the appropriate nucleus calls. The <i>dcFree</i> nucleus call destroys the data cache object specified by the <i>lcIdp</i> argument. The External Proxy-Mapper can destroy a data cache object only if it is empty (it does not contain any physical pages) and the corresponding data segment is used neither by the Proxy-Mapper nor by the VM.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	If successful <code>K_OK</code> is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	[ <code>K_ENOMEM</code> ]                      The system is out of resources.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>dcFlush(2SEG)</i> , <i>dcRead(2SEG)</i> , <i>dcWrite(2SEG)</i> , <i>dcFillZero(2SEG)</i> , <i>dcPgNumber(2SEG)</i> , <i>dcCluster(2SEG)</i> , <i>dcSync(2SEG)</i>				

```
, dcPxmDeclare(2SEG) , rgnMapFromDtCache(2SEG) ,  
rgnInitFromDtCache(2SEG) , rgnFlush(2SEG)
```

<b>NAME</b>	dcCluster – Set the input and output cluster sizes of a data cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcCluster(KnLcId lcId, VmSize incluster, VmSize outcluster);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcCluster</i> function sets the cluster size attributes of the data cache specified by the <i>lcId</i> argument. The <i>incluster</i> argument specifies the preferred size for both the <i>pullIn</i> and <i>getAcc</i> up-calls. These invocations will therefore try to use page lists starting on a boundary multiple of <i>incluster</i> and of a size equal to <i>incluster</i> (<i>there is no guarantee it will always be the case</i>).</p> <p>Similarly, the <i>outcluster</i> argument specifies the preferred size to be used by the nucleus for building page lists to be passed to the proxy-mapper during the <i>pushOutAsyn</i> up-call.</p> <p>The default size in both cases is the page size. The nucleus takes the cluster size into account if it is a power of two and greater than the page size. This call does not acquire any VM lock and therefore cannot be blocked.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	If successful this call returns K_OK, otherwise the appropriate error code is returned.				
<b>ERRORS</b>	[K_EROUND]                      At least one of the sizes is not a multiple of the page size.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>dcAlloc(2SEG)</code> , <code>dcFree(2SEG)</code> , <code>dcFlush(2SEG)</code> , <code>dcFillZero(2SEG)</code> , <code>dcRead(2SEG)</code> , <code>dcWrite(2SEG)</code> , <code>dcPgNumber(2SEG)</code> , <code>dcSync(2SEG)</code> , <code>dcPxmDeclare(2SEG)</code> , <code>rgnMapFromDtCache(2SEG)</code> , <code>rgnInitFromDtCache(2SEG)</code> , <code>rgnFlush(2SEG)</code>				

<b>NAME</b>	dcFillZero – Fill a data segment with zero				
<b>SYNOPSIS</b>	#include <mem/chMem.h> int dcFillZero(KnLcid <i>lcid</i> , VmOffset <i>start</i> , VmOffset <i>end</i> , VmOffset <i>zeroStart</i> );				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.  The <i>dcFillZero</i> function gives write access rights to the nucleus for the range of bytes between <i>start</i> and <i>end</i> . The <i>start</i> and <i>end+1</i> arguments must be “fragment”-aligned. A fragment is defined as the unit used by the VM to manage access rights for a segment. The size of a fragment is usually set to 512 bytes but may be implementation-dependent. It also fills with zeros the range of bytes between the <i>startZero</i> and <i>end</i> arguments. The <i>startZero</i> argument has no alignment constraints, however, it must be within the range of bytes defined by the <i>start</i> and the <i>end</i> arguments. The corresponding part of the data cache is marked as modified.				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	[K_EROUND]                      At least one of the arguments is not “fragment”-aligned.  [K_EOFFSET]                      The <i>startZero</i> offset is not within the range defined by <i>start</i> , <i>end</i> .				
<b>ATTRIBUTES</b>	See attributes(5) for descriptions of the following attributes: <table border="1" style="margin-left: 20px; width: 100%;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	dcAlloc(2SEG), dcFlush(2SEG), dcFree(2SEG), dcRead(2SEG), dcWrite(2SEG), dcPgNumber(2SEG), dcCluster(2SEG), dcSync(2SEG), dcPxmDeclare(2SEG), rgnMapFromDtCache(2SEG), rgnInitFromDtCache(2SEG), rgnFlush(2SEG)				



<b>NAME</b>	dcSync, dcFlush – Sync a data cache object; Flush a data cache object
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcSync(KnLcid lcid, int pagenb, VmOffset * offset);  int dcFlush(KnLcid lcid, KnPxmFlushReq * fldesc, void * pout);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcSync</i> function scans the physical pages attached to the data cache object specified by the <i>lcid</i> argument. The scanning is performed incrementally by offset, the start page is defined by the <i>offset</i> argument. The <i>pagenb</i> defines the maximum number of pages to be scanned. Each modified page will be passed to the <i>pushOutAsyn</i> up-call. The way the pages are passed to this upcall depends on the out cluster size as defined by <i>dcCluster</i>. The function returns either a non-negative value, which means that all pages have been scanned. The return value is the difference between <i>pagenb</i> and the number of scanned pages. A value of -1 means that all pages specified have been scanned but that there are still unscanned pages. In such a case, <i>dcSync</i> sets the <i>offset</i> output argument to the first unscanned page offset.</p> <p>The <i>dcFlush</i> nucleus call flushes physical pages associated to the data cache object defined by the <i>lcid</i> argument, according to the structure pointed to by <i>fldesc</i>. This structure is composed of the following members:</p> <pre>VmFlags    flags ; VmOffset   desStart ; VmOffset   desEnd ; VmOffset   reqStart ; VmOffset   reqEnd ;</pre> <p>The <i>dcFlush</i> call enables the flushing of all pages in the range defined by the <i>reqStart</i> and <i>reqEnd</i> offsets to be requested. It also permits the specification of a desired range of pages to be flushed, by setting the <i>desStart</i> and <i>desEnd</i> fields to the appropriate values. If no desired range is needed, the <i>desStart</i> field should be set to the value of the <i>reqStart</i> field. Similarly, the <i>desEnd</i> field should be set the value of the <i>reqEnd</i> field. The <i>desStart</i> argument should be less than or equal to <i>reqStart</i>, and the <i>desEnd</i> argument should be greater than or equal to <i>reqEnd</i>. The system will flush all pages in the required range, and will try to flush pages in the desired range. In particular, it will avoid being blocked for pages which are in the desired range but not in the required range.</p> <p>The <i>dcFlush</i> call may perform the flush in different ways according to the value of the <i>flags</i> field. This field specifies the required flush mode and consists of several parts:</p>

K_FLUSH_ACCREC_NONE	Access rights are left unchanged. This is useful for flushing dirty pages without modifying access rights previously granted to the data cache object.
K_FLUSH_ACCREC_WRITE	Write access rights are recalled.
K_FLUSH_ACCREC_READ	Read and write access rights are recalled.
K_FLUSH_ACCREC_INVALID	Read and write access rights are recalled and pages in the desired range are destroyed.
K_FLUSH_POUT_NONE	The <i>pushOutAsyn</i> up-call is never invoked. This is useful for modifying access rights without performing any push out operations.
K_FLUSH_POUT_DIRTY	The <i>pushOutAsyn</i> up-call is performed for dirty pages only.
K_FLUSH_POUT_ALL	The <i>pushOutAsyn</i> up-call is invoked for all pages even if they are not dirty.
K_FLUSH_PAGEFAULT	If this flag is set, <i>dcFlush</i> sets the K_PAGEFAULT flag in the <i>flags</i> argument of the <i>pushOutAsyn</i> up-call.
K_FLUSH_PAGELIST	If this flag is set, <i>dcFlush</i> will invoke the <i>pushOutAsyn</i> up-call with lists of contiguous pages not aligned on the cluster size defined for the data cache object. If possible, the entire desired range will be passed as a single list to the <i>pushOutAsyn</i> up-call. Several lists will be built in the case of discontinuous lists of pages (in such a case, several up-calls will be invoked).

The *pout* argument is an opaque for the VM and will be passed back as an argument of the *pushOutAsyn* up-call, if any.

If successful, *dcFlush* returns *K\_OK*, otherwise an error code is returned.

The *dcFlush* call acquires an exclusive lock on pages being flushed. For pages belonging to the desired range but not to the required range, *dcFlush* does not block in order to acquire this lock; if the lock cannot be acquired the page will not be flushed. Pages specified in the *getAcc* up-call are not visible to the *dcFlush* system call. Thus, VM insures that the *dcFlush* operation will never be blocked by a *getAcc* up-call. In other words, there is no risk of deadlock (from a VM point of view) when an Extrenal Proxy-Mapper performs a *dcFlush* while a *getAcc* up-call is blocked in this PXM.

**RESTRICTIONS**

The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE**

If successful *dcSync* returns -1 when all pages specified have been scanned, but there are still unscanned pages. Otherwise, it returns the difference between the number of scanned pages and the *pageNb* argument.

The *dcFlush* call returns *K\_OK* in case of success, otherwise an error code is returned.

**ERRORS**

[K\_EROUND] At least one of the sizes is not a multiple of the page size.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*dcAlloc(2SEG)* , *dcFree(2SEG)* , *dcFlush(2SEG)* , *dcFillZero(2SEG)* , *dcRead(2SEG)* , *dcWrite(2SEG)* , *dcPgNumber(2SEG)* , *dcSync(2SEG)* , *dcPxmDeclare(2SEG)* , *rgnMapFromDtCache(2SEG)* , *rgnInitFromDtCache(2SEG)* , *rgnFlush(2SEG)*

**NAME** | dcAlloc, dcFree – Allocate a data cache for a segment; Free a previously allocated data cache

**SYNOPSIS** | #include <mem/chMem.h>  
 int **dcAlloc**(KnSgId *sgid*, KnExtPxMapper \* *pxm*, VmFlags *flags*, KnLcId \* *lcidp*);  
 void **dcFree**(KnLcId *lcid*);

**FEATURES** | PXM\_EXT

**DESCRIPTION** | **Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.  
 The *dcAlloc* function allocates a new data cache object associated with the data segment identified by the *sgId* parameter. Two subsequent calls to *dcAlloc* with the same segment identifier will allocate two different data cache objects. It is thus the responsibility of the invoker to insure a one-to- one mapping between data cache objects and data segments. The *sgId* is an opaque value for the nucleus which is only used as an argument within upcalls from the VM to the external Proxy-Mapper managing the data segment. The *pxm* argument specifies the routines which will be used to perform upcalls required for this data segment (see *dcPxmDeclare(2SEG)* ). The *flags* argument specifies the required properties of the data segment. It must be either *K\_NOSWAPOUT* if the data cache must not be swapped, or 0. If successful, *dcAlloc* returns *K\_OK* and sets the *lcidp* output argument to the identifier of the newly allocated data cache object, otherwise an error code is returned and the value of the *lcidp* field is undefined. The *lcidp* returned by *dcAlloc* is an opaque for the External Proxy-Mapper and must only be used as an argument for the appropriate nucleus calls. The *dcFree* nucleus call destroys the data cache object specified by the *lcidp* argument. The External Proxy-Mapper can destroy a data cache object only if it is empty (it does not contain any physical pages) and the corresponding data segment is used neither by the Proxy-Mapper nor by the VM.

**RESTRICTIONS** | The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE** | If successful *K\_OK* is returned, otherwise a negative error code is returned.

**ERRORS** | [*K\_ENOMEM*] The system is out of resources.

**ATTRIBUTES** | See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | *dcFlush(2SEG)* , *dcRead(2SEG)* , *dcWrite(2SEG)* , *dcFillZero(2SEG)* , *dcPgNumber(2SEG)* , *dcCluster(2SEG)* , *dcSync(2SEG)*

```
, dcPxmDeclare(2SEG) , rgnMapFromDtCache(2SEG) ,  
rgnInitFromDtCache(2SEG) , rgnFlush(2SEG)
```

<b>NAME</b>	dcGetPages – Get a list of pages for read-ahead purpose
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcGetPages(KnLcid lcid, VmOffset start, VmOffset end, VmFlags flags, KnPage **page);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcGetPages</i> function enables a proxy-mapper to implement a read-ahead policy. It scans the data cache identified by the <i>lcid</i> argument, looking for pages belonging to the range defined by the <i>start</i> and <i>end</i> arguments. A new physical page is allocated for each page missing in the data cache. These pages are grouped in a contiguous list whose head is returned to the invoker at the location defined by the argument <i>page</i>. They may then be mapped or unmapped using the <i>pageMap</i> and <i>pageUnmap</i> nucleus calls, respectively. When pages have been filled with meaningful data, the proxy-mapper informs the nucleus using the <i>pageIoDone</i> nucleus call.</p> <p>When I/O is completed, access rights are granted to the nucleus according to the values of the <i>flags</i> argument. If <i>K_READABLE</i> is set, read access is granted to the nucleus. If <i>K_WRITABLE</i> is set write access is granted to the nucleus. Pages returned by the nucleus to the proxy-mapper as a result of the <i>dcGetPages</i> nucleus call are write-locked until the read-ahead completes. This insures that concurrent access to these pages (mapping, reading, flushing,...) will be blocked until the read-ahead completes.</p> <p>If the <i>K_NOWAITFORMEMORY</i> value is set in the <i>flags</i> argument, the nucleus will not allocate free pages in case free memory is low.</p> <p>The page list returned by the <i>dcGetPages</i> nucleus call starts with the first missing page in the cache and will end at the first non missing page. However, this list cannot be greater than the range specified by the <i>start</i> and <i>end</i> arguments (which must be page aligned). The effective size of the list is also constrained by the input cluster size of the data cache, except if the <i>K_PAGELIST</i> value is set within the <i>flags</i> argument.</p>
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.
<b>ERRORS</b>	[K_EROUND]                      One of the <i>start</i> or <i>end</i> arguments is not page aligned.
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

dcAlloc(2SEG), dcFlush(2SEG), dcFree(2SEG), dcRead(2SEG),  
dcWrite(2SEG), dcPgNumber(2SEG), dcCluster(2SEG),  
dcSync(2SEG), dcPxmDeclare(2SEG), rgnMapFromDtCache(2SEG),  
rgnInitFromDtCache(2SEG), rgnFlush(2SEG)

<b>NAME</b>	dcIsDirty – Test and reset data cache dirty bit				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcIsDirty(KnLcId lcid);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcIsDirty</i> function returns 0 if the data cache defined by the <i>lcid</i> argument has not been modified through mapping. It returns a non-null value if the data cache defined by the <i>lcid</i> argument has been dirtied through mapping. The data cache dirty bit is set whenever a dirty mapped page is unloaded from mmu tables or is sync'ed. This may happen when a <i>rgnFree</i> nucleus call occurs or when a <i>dcSync</i> or <i>dcFlush</i> call is invoked. After the dirty bit of the data cache has been tested, it is reset. Thus, a subsequent call to <i>dcIsDirty</i> will return a null value. Modification of a data cache through a call to <i>dcWrite</i> will not set the data cache dirty bit.</p> <p>This call is provided mainly to enable filesystems to determine whether a file has been modified through mapping.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	This function returns 0 if the data cache is not dirty, 1 otherwise.				
<b>ERRORS</b>	No error messages are returned.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>dcAlloc(2SEG)</i> , <i>dcFlush(2SEG)</i> , <i>dcFree(2SEG)</i> , <i>dcRead(2SEG)</i> , <i>dcWrite(2SEG)</i> , <i>dcPgNumber(2SEG)</i> , <i>dcCluster(2SEG)</i> , <i>dcSync(2SEG)</i> , <i>dcPxmDeclare(2SEG)</i> , <i>rgnMapFromDtCache(2SEG)</i> , <i>rgnInitFromDtCache(2SEG)</i> , <i>rgnFlush(2SEG)</i>				



<b>NAME</b>	dcPgNumber – Get the number of pages for a data cache object				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcPgNumber(KnLcId lcid);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcPgNumber</i> function returns the number of pages attached to the data cache object specified by the <i>lcid</i> argument. This argument should have been acquired via a previous call to <i>dcAlloc</i>. This call does not acquire a VM lock and therefore cannot be blocked.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	This call returns the number of pages.				
<b>ERRORS</b>	No error messages are returned.				
<b>ATTRIBUTES</b>	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<code>dcAlloc(2SEG)</code> , <code>dcFree(2SEG)</code> , <code>dcFlush(2SEG)</code> , <code>dcFillZero(2SEG)</code> , <code>dcRead(2SEG)</code> , <code>dcWrite(2SEG)</code> , <code>dcCluster(2SEG)</code> , <code>dcSync(2SEG)</code> , <code>dcPxmDeclare(2SEG)</code> , <code>rgnInitFromDtCache(2SEG)</code> , <code>rgnFlush(2SEG)</code>				

**NAME** | dcPxmDeclare – Initialize an external proxy-mapper descriptor

**SYNOPSIS** | #include <mem/chMem.h>  
 int dcPxmDeclare(KnExtPxMapper \*pxm);

**FEATURES** | PXM\_EXT

**DESCRIPTION** | **Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.

The *dcPxmDeclare* function completes the initialization of the proxy-mapper descriptor pointed to by the *pxm* argument. This descriptor can be used later within *dcAlloc* nucleus calls. The proxy-mapper descriptor should have previously been initialized by the invoker. This is done by creating the appropriate laps to fill the following fields of the *KnExtPxMapper* data structure:

```
KnLapDesc  openLap ;
KnLapDesc  closeLap ;
KnLapDesc  getAccLap ;
KnLapDesc  pullInLap ;
KnLapDesc  pushOutAsynLap ;
KnLapDesc  swapOutLap ;
KnLapDesc  statLap ;
```

All local access point descriptors must be valid and remain valid as long as there are data caches managed by a proxy-mapper. As the proxy-mapper descriptor is not copied within the nucleus address space, this descriptor must remain valid and accessible as long as there are data caches managed by a proxy-mapper.

Each of the laps correspond to a possible upcall made by the nucleus to perform actions on the data segments managed by that proxy-mapper. The syntax and semantics of these upcalls are described in the appropriate manual pages.

The *swapOutLap* may be initialised to 0, usually by means of the *lapDescZero* nucleus call. This will prevent the nucleus from performing this upcall at swapout time.

**RESTRICTIONS** | The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE** | K\_OK is always returned.

**ERRORS** | No error messages are returned.

**ATTRIBUTES** | See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | *dcAlloc(2SEG)*, *svLapCreate(2K)*, *lapDescZero(2K)*

<b>NAME</b>	dcRead, dcWrite – Read data from a data cache; Write data to a data cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcRead(KnLcId <i>lcid</i>, VmOffset <i>start</i>, VmOffset <i>end</i>, KnCap * <i>actcap</i>, VmAddr <i>addr</i>, void * <i>acclck</i>);  int dcWrite(KnLcId <i>lcid</i>, VmOffset <i>start</i>, VmOffset <i>end</i>, KnCap * <i>actcap</i>, VmAddr <i>addr</i>, void * <i>acclck</i>);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcRead</i> function reads data from the data cache object named by <i>lcid</i> into an actor address space. The source data is specified by the <i>start</i> and <i>end</i> offsets which define the offset of the first byte to be read and the offset of the last byte to be read. These offsets are related to the data segment associated with the data cache object of a previous <i>dcAlloc</i>. The data is copied, starting at the location defined by the <i>addr</i> input/output argument, within the actor defined by the <i>actcap</i> argument. If <i>actcap</i> is set to the value <code>K_MYACTOR</code> the destination actor is the current one. The <i>addr</i> argument is updated by the <i>dcRead</i> nucleus call, and upon return, points to an address beyond the last read byte upon return. The <i>acclck</i> argument is an opaque field for the nucleus, which will be passed to the external Proxy-Mapper as an argument of the <i>getAcc</i> up-call, if any.</p> <p>The <i>dcWrite</i> function writes data from an actor address space to the data cache object named by <i>lcid</i>. The destination data is specified by the <i>start</i> and <i>end</i> offsets which define the offset of the first byte to be written and the offset of the last byte to be written. These offsets are related to the data segment associated with the data cache object by a previous <i>dcAlloc</i>. The data is copied, starting at the location defined by the <i>addr</i> input/output argument, within the actor defined by the <i>actcap</i> argument. If <i>actcap</i> is set to the value <code>K_MYACTOR</code> the source actor is the current one. The <i>addr</i> argument is updated by the <i>dcWrite</i> nucleus call, and upon return, points to an address beyond the last byte written. The <i>acclck</i> argument is an opaque field for the nucleus, which will be passed to the external Proxy-Mapper as an argument of the <i>getAcc</i> up-call, if any.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	If successful <code>K_OK</code> is returned, otherwise a negative error code is returned. In either case, the <i>addr</i> argument is set to an address beyond the last byte read/written.				
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EADDR]</td> <td>The address is outside any allocated region.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's or target's address space.</td> </tr> </table>	[K_EADDR]	The address is outside any allocated region.	[K_EFAULT]	Some of the arguments provided are outside the caller's or target's address space.
[K_EADDR]	The address is outside any allocated region.				
[K_EFAULT]	Some of the arguments provided are outside the caller's or target's address space.				

- [K\_EINVAL] An inconsistent actor capability was given.
- [K\_EOFFSET] Attempt to access a segment outside its valid offset range.
- [K\_EPROT] Attempt to write a read only data cache.
- [K\_EUNKNOWN] *actcap* does not specify a reachable actor.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`dcAlloc(2SEG)`, `dcFree(2SEG)`, `dcFlush(2SEG)`, `dcFillZero(2SEG)`  
`, dcPgNumber(2SEG)`, `dcCluster(2SEG)`, `dcSync(2SEG)`  
`, dcPxmDeclare(2SEG)`, `rgnMapFromDtCache(2SEG)`,  
`rgnInitFromDtCache(2SEG)`, `rgnFlush(2SEG)`

<b>NAME</b>	dcSync, dcFlush – Sync a data cache object; Flush a data cache object
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcSync(KnLcid lcid, int pagenb, VmOffset * offset);  int dcFlush(KnLcid lcid, KnPxmFlushReq * fldesc, void * pout);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcSync</i> function scans the physical pages attached to the data cache object specified by the <i>lcid</i> argument. The scanning is performed incrementally by offset, the start page is defined by the <i>offset</i> argument. The <i>pagenb</i> defines the maximum number of pages to be scanned. Each modified page will be passed to the <i>pushOutAsyn</i> up-call. The way the pages are passed to this upcall depends on the out cluster size as defined by <i>dcCluster</i>. The function returns either a non-negative value, which means that all pages have been scanned. The return value is the difference between <i>pagenb</i> and the number of scanned pages. A value of -1 means that all pages specified have been scanned but that there are still unscanned pages. In such a case, <i>dcSync</i> sets the <i>offset</i> output argument to the first unscanned page offset.</p> <p>The <i>dcFlush</i> nucleus call flushes physical pages associated to the data cache object defined by the <i>lcid</i> argument, according to the structure pointed to by <i>fldesc</i>. This structure is composed of the following members:</p> <pre>VmFlags    flags ; VmOffset   desStart ; VmOffset   desEnd ; VmOffset   reqStart ; VmOffset   reqEnd ;</pre> <p>The <i>dcFlush</i> call enables the flushing of all pages in the range defined by the <i>reqStart</i> and <i>reqEnd</i> offsets to be requested. It also permits the specification of a desired range of pages to be flushed, by setting the <i>desStart</i> and <i>desEnd</i> fields to the appropriate values. If no desired range is needed, the <i>desStart</i> field should be set to the value of the <i>reqStart</i> field. Similarly, the <i>desEnd</i> field should be set to the value of the <i>reqEnd</i> field. The <i>desStart</i> argument should be less than or equal to <i>reqStart</i>, and the <i>desEnd</i> argument should be greater than or equal to <i>reqEnd</i>. The system will flush all pages in the required range, and will try to flush pages in the desired range. In particular, it will avoid being blocked for pages which are in the desired range but not in the required range.</p> <p>The <i>dcFlush</i> call may perform the flush in different ways according to the value of the <i>flags</i> field. This field specifies the required flush mode and consists of several parts:</p>

K_FLUSH_ACCREC_NONE	Access rights are left unchanged. This is useful for flushing dirty pages without modifying access rights previously granted to the data cache object.
K_FLUSH_ACCREC_WRITE	Write access rights are recalled.
K_FLUSH_ACCREC_READ	Read and write access rights are recalled.
K_FLUSH_ACCREC_INVALID	Read and write access rights are recalled and pages in the desired range are destroyed.
K_FLUSH_POUT_NONE	The <i>pushOutAsyn</i> up-call is never invoked. This is useful for modifying access rights without performing any push out operations.
K_FLUSH_POUT_DIRTY	The <i>pushOutAsyn</i> up-call is performed for dirty pages only.
K_FLUSH_POUT_ALL	The <i>pushOutAsyn</i> up-call is invoked for all pages even if they are not dirty.
K_FLUSH_PAGEFAULT	If this flag is set, <i>dcFlush</i> sets the K_PAGEFAULT flag in the <i>flags</i> argument of the <i>pushOutAsyn</i> up-call.
K_FLUSH_PAGELIST	If this flag is set, <i>dcFlush</i> will invoke the <i>pushOutAsyn</i> up-call with lists of contiguous pages not aligned on the cluster size defined for the data cache object. If possible, the entire desired range will be passed as a single list to the <i>pushOutAsyn</i> up-call. Several lists will be built in the case of discontinuous lists of pages (in such a case, several up-calls will be invoked).

The *pout* argument is an opaque for the VM and will be passed back as an argument of the *pushOutAsyn* up-call, if any.

If successful, *dcFlush* returns *K\_OK*, otherwise an error code is returned.

The *dcFlush* call acquires an exclusive lock on pages being flushed. For pages belonging to the desired range but not to the required range, *dcFlush* does not block in order to acquire this lock; if the lock cannot be acquired the page will not be flushed. Pages specified in the *getAcc* up-call are not visible to the *dcFlush* system call. Thus, VM insures that the *dcFlush* operation will never be blocked by a *getAcc* up-call. In other words, there is no risk of deadlock (from a VM point of view) when an Extrenal Proxy-Mapper performs a *dcFlush* while a *getAcc* up-call is blocked in this PXM.

**RESTRICTIONS**

The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE**

If successful *dcSync* returns -1 when all pages specified have been scanned, but there are still unscanned pages. Otherwise, it returns the difference between the number of scanned pages and the *pageNb* argument.

The *dcFlush* call returns *K\_OK* in case of success, otherwise an error code is returned.

**ERRORS**

[K\_EROUND]                      At least one of the sizes is not a multiple of the page size.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*dcAlloc(2SEG)* , *dcFree(2SEG)* , *dcFlush(2SEG)* , *dcFillZero(2SEG)* , *dcRead(2SEG)* , *dcWrite(2SEG)* , *dcPgNumber(2SEG)* , *dcSync(2SEG)* , *dcPxmDeclare(2SEG)* , *rgnMapFromDtCache(2SEG)* , *rgnInitFromDtCache(2SEG)* , *rgnFlush(2SEG)*

**NAME** | dcTrunc - Truncate a data segment

**SYNOPSIS** | #include <mem/chMem.h>  
 int dcTrunc(KnLcId lcid, VmOffset start, VmOffset end, VmOffset zeroStart);

**FEATURES** | PXM\_EXT

**DESCRIPTION** | **Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.  
 The *dcTrunc* call grants write access to the nucleus for the range of bytes between *start* and *end*, and invalidates the [*end+1*, *K\_MAXVMOFFSET*] range of the data cache specified by the *lcid* argument. It then fills with zeros the range of bytes between the *zeroStart* and *end* arguments, and marks the range of bytes defined by the *start*, *end* arguments as modified. The *start* and *end* arguments must be fragment-aligned, and the *zeroStart* argument must be included between these two limits.  
 Logically, *dcTrunc* is a combination of *dcFlush* and *dcFillZero*. The VM guarantees that the *dcTrunc* operation will never be blocked by a *getAcc* up-call.

**RESTRICTIONS** | The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE** | If successful K\_OK is returned, otherwise a negative error code is returned.

**ERRORS** | [K\_EROUND] | At least one of the arguments is not fragment-aligned.  
 [K\_EOFFSET] | The *startZero* offset is not within the range defined by *start*, *end*.

**ATTRIBUTES** | See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | dcAlloc(2SEG), dcFlush(2SEG), dcFillZero(2SEG), dcFree(2SEG), dcRead(2SEG), dcWrite(2SEG), dcPgNumber(2SEG), dcCluster(2SEG), dcSync(2SEG), dcPxmDeclare(2SEG), rgnMapFromDtCache(2SEG), rgnInitFromDtCache(2SEG), rgnFlush(2SEG)



<b>NAME</b>	dcRead, dcWrite – Read data from a data cache; Write data to a data cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int dcRead(KnLcId <i>lcid</i>, VmOffset <i>start</i>, VmOffset <i>end</i>, KnCap * <i>actcap</i>, VmAddr <i>addr</i>, void * <i>acclck</i>);  int dcWrite(KnLcId <i>lcid</i>, VmOffset <i>start</i>, VmOffset <i>end</i>, KnCap * <i>actcap</i>, VmAddr <i>addr</i>, void * <i>acclck</i>);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>dcRead</i> function reads data from the data cache object named by <i>lcid</i> into an actor address space. The source data is specified by the <i>start</i> and <i>end</i> offsets which define the offset of the first byte to be read and the offset of the last byte to be read. These offsets are related to the data segment associated with the data cache object of a previous <i>dcAlloc</i>. The data is copied, starting at the location defined by the <i>addr</i> input/output argument, within the actor defined by the <i>actcap</i> argument. If <i>actcap</i> is set to the value <code>K_MYACTOR</code> the destination actor is the current one. The <i>addr</i> argument is updated by the <i>dcRead</i> nucleus call, and upon return, points to an address beyond the last read byte upon return. The <i>acclck</i> argument is an opaque field for the nucleus, which will be passed to the external Proxy-Mapper as an argument of the <i>getAcc</i> up-call, if any.</p> <p>The <i>dcWrite</i> function writes data from an actor address space to the data cache object named by <i>lcid</i>. The destination data is specified by the <i>start</i> and <i>end</i> offsets which define the offset of the first byte to be written and the offset of the last byte to be written. These offsets are related to the data segment associated with the data cache object by a previous <i>dcAlloc</i>. The data is copied, starting at the location defined by the <i>addr</i> input/output argument, within the actor defined by the <i>actcap</i> argument. If <i>actcap</i> is set to the value <code>K_MYACTOR</code> the source actor is the current one. The <i>addr</i> argument is updated by the <i>dcWrite</i> nucleus call, and upon return, points to an address beyond the last byte written. The <i>acclck</i> argument is an opaque field for the nucleus, which will be passed to the external Proxy-Mapper as an argument of the <i>getAcc</i> up-call, if any.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	If successful <code>K_OK</code> is returned, otherwise a negative error code is returned. In either case, the <i>addr</i> argument is set to an address beyond the last byte read/written.				
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EADDR]</td> <td>The address is outside any allocated region.</td> </tr> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's or target's address space.</td> </tr> </table>	[K_EADDR]	The address is outside any allocated region.	[K_EFAULT]	Some of the arguments provided are outside the caller's or target's address space.
[K_EADDR]	The address is outside any allocated region.				
[K_EFAULT]	Some of the arguments provided are outside the caller's or target's address space.				

- [K\_EINVAL] An inconsistent actor capability was given.
- [K\_EOFFSET] Attempt to access a segment outside its valid offset range.
- [K\_EPROT] Attempt to write a read only data cache.
- [K\_EUNKNOWN] *actcap* does not specify a reachable actor.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`dcAlloc(2SEG)`, `dcFree(2SEG)`, `dcFlush(2SEG)`, `dcFillZero(2SEG)`  
`dcPgNumber(2SEG)`, `dcCluster(2SEG)`, `dcSync(2SEG)`  
`dcPxmDeclare(2SEG)`, `rgnMapFromDtCache(2SEG)`,  
`rgnInitFromDtCache(2SEG)`, `rgnFlush(2SEG)`

<b>NAME</b>	lcOpen, lcClose, lcCap – Find or create a local cache object for a segment; Release a local cache object; Return the capability of a local cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcOpen(KnCap * sgcap, VmFlags flags, int * cachelip);  int lcCap(int cacheli, KnCap * lccap);  int lcClose(int cacheli);</pre>				
<b>FEATURES</b>	MEM_VM				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcOpen</i> call allows the caller to force the kernel to use the same local cache object for the segment specified until a subsequent <i>lcClose</i> call is received.</p> <p>The <i>sgcap</i> argument is a pointer to the target segment capability. The <i>flags</i> argument must be zero. This argument will be used in future extensions of the interface.</p> <p>The <i>lcOpen</i> call returns into the variable pointed to by the <i>cachelip</i> argument, a local id of the local cache object corresponding to the target segment. The caller can then use the local id in a number of subsequent <i>lcRead</i>, <i>lcWrite</i> (see <i>lcRead(2SEG)</i>) and <i>lcCap</i> (see below) calls. The local id must be released using the <i>lcClose</i> call and should not be used any more.</p> <p>The <i>lcCap</i> call finds the local cache object specified by its local id ( <i>cacheli</i> argument) and returns the local cache capability in the variable pointed to by the <i>lccap</i> argument.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors. An attempt to pass an invalid local id may produce unpredictable system behavior.				
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EFAULT]                      Some of the arguments provided are outside the caller's address space.</p> <p>[K_ENOMEM]                      The system is out of resources.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>vmStat(2K)</i> , <i>MpGetAccess(2SEG)</i> , <i>MpPullIn(2SEG)</i> , <i>rgnMap(2SEG)</i> , <i>lcFlush(2SEG)</i> , <i>lcFillZero(2SEG)</i> , <i>lcTrunc(2SEG)</i>				

<b>NAME</b>	lcOpen, lcClose, lcCap – Find or create a local cache object for a segment; Release a local cache object; Return the capability of a local cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcOpen(KnCap * sgcap, VmFlags flags, int * cachelip);  int lcCap(int cacheli, KnCap * lccap);  int lcClose(int cacheli);</pre>				
<b>FEATURES</b>	MEM_VM				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcOpen</i> call allows the caller to force the kernel to use the same local cache object for the segment specified until a subsequent <i>lcClose</i> call is received.</p> <p>The <i>sgcap</i> argument is a pointer to the target segment capability. The <i>flags</i> argument must be zero. This argument will be used in future extensions of the interface.</p> <p>The <i>lcOpen</i> call returns into the variable pointed to by the <i>cachelip</i> argument, a local id of the local cache object corresponding to the target segment. The caller can then use the local id in a number of subsequent <i>lcRead</i>, <i>lcWrite</i> (see <i>lcRead(2SEG)</i>) and <i>lcCap</i> (see below) calls. The local id must be released using the <i>lcClose</i> call and should not be used any more.</p> <p>The <i>lcCap</i> call finds the local cache object specified by its local id ( <i>cacheli</i> argument) and returns the local cache capability in the variable pointed to by the <i>lccap</i> argument.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors. An attempt to pass an invalid local id may produce unpredictable system behavior.				
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	<table border="0"> <tr> <td style="padding-right: 20px;">[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>The system is out of resources.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_ENOMEM]	The system is out of resources.
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.				
[K_ENOMEM]	The system is out of resources.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>vmStat(2K)</i> , <i>MpGetAccess(2SEG)</i> , <i>MpPullIn(2SEG)</i> , <i>rgnMap(2SEG)</i> , <i>lcFlush(2SEG)</i> , <i>lcFillZero(2SEG)</i> , <i>lcTrunc(2SEG)</i>				

<b>NAME</b>	lcFillZero – zero a range of a local cache
<b>SYNOPSIS</b>	#include <mem/chMem.h> int lcFillZero(KnObjDesc *lcdesc, VmFlags flags, unsigned longordernb);
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcFillZero</i> call grants write access rights and zero-fills and marks as modified the [<i>lcdesc-&gt;startOffset</i>, <i>lcdesc-&gt;startOffset</i> + <i>lcdesc-&gt;size</i>) range of the local cache specified by the <i>lcdesc-&gt;dataObject</i> capability. The <i>KnObjDesc</i> structure is described in <i>lcFlush(2SEG)</i>.</p> <p>The <i>lcdesc-&gt;startOffset</i> and <i>lcdesc-&gt;size</i> fields must be fragment-aligned. The size of the fragment is implementation-dependent and is usually equal to the virtual page size divided by 8 (number of bits in one byte).</p> <p>The <i>flags</i> argument must be zero and will be used in future interface extensions. The <i>ordernb</i> argument is ignored.</p> <p>The <i>lcFillZero</i> operation will never be blocked either by a pure <i>MpGetAccess</i> or by an impure <i>MpPullIn</i> up-call.</p>
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.
<b>ERRORS</b>	<p>[K_EFAULT]                   Some of the arguments provided are outside the caller's address space.</p> <p>[K_EINVAL]                   An inconsistent local cache capability was provided.</p> <p>[K_EUNKNOWN]                <i>lcdesc-&gt;dataObject</i> does not specify a reachable local cache.</p> <p>[K_EROUND]                   <i>lcdesc-&gt;startOffset</i> or <i>lcdesc-&gt;size</i> is not fragment-aligned.</p> <p>[K_EOFFSET]                 Tried to fill a segment outside the valid offset range in a segment, as returned by <i>vmStat</i>.</p> <p>[K_EINVAL]                   The <i>flags</i> argument contains invalid flag values.</p> <p>[K_EFAIL]                    An <i>ipcCall</i> transaction failed during the remote <i>lcFillZero</i>.</p> <p>[K_EMAPPER]                 The mapper doesn't respect the vm/mapper protocol.</p>
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

MpGetAccess(2SEG), MpPullIn(2SEG), rgnMap(2SEG), lcOpen(2SEG),  
lcFlush(2SEG), lcTrunc(2SEG)

<b>NAME</b>	lcFlush, lcSetRights, sgFlush, sgSyncAll, vmFlush – flush local cache(s)
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcFlush(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int lcSetRights(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int sgFlush(KnObjDesc * segdesc, VmFlags flags);  int sgSyncAll(void);  int vmFlush(KnCap * actorcap, VmAddr address, VmSize size, VmFlags flags);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcFlush</i> call performs a flush operation on a range of a local cache. The <i>lcdesc</i> argument points to a <i>KnObjDesc</i> structure whose members are the following:</p> <pre>KnCap      dataObject ; VmOffset   startOffset ; VmSize     size ;</pre> <p>where the <i>dataObject</i> field is the local cache capability, the <i>startOffset</i> field is the range start offset in the local cache, and the <i>size</i> field is the range size. Both <i>startOffset</i> and <i>size</i> must be fragment-aligned. The size of the fragment is implementation-dependent and is usually equal to the virtual page size divided by 8 (number of bits in one byte).</p> <p>For <i>lcFlush</i>, the target local cache is directly specified by its capability; the flush operation is applied even if the cache is managed by a remote site.</p> <p>The <i>sgFlush</i> call takes the same arguments, except the <i>dataObject</i> field specifies the capability of the cached segment. For <i>sgFlush</i>, the target local cache is indirectly specified by the capability of the corresponding segment; the flush operation implicitly applies to the segment's cache(s) located on the caller's site.</p> <p>The <i>vmFlush</i> call performs a flush operation on the ranges of the local caches mapped to an address range of an actor address space. The address range must be page-aligned.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <i>K_MYACTOR</i>, the address space of the current actor is used. If <i>actorcap</i> is <i>K_SVACTOR</i>, the supervisor address space is used.</p> <p>The <i>flags</i> argument specifies the mode of the flush and the upper access rights for the target parts after the flush operation.</p>

If the `K_COPYBACK` flag is set, *xxFlush* writes back any (even clean) cached data of the target parts and keeps the parts access rights unchanged.

If the `K_WRITABLE` flag is set, *xxFlush* writes back all modified data of the target parts and keeps the parts access rights unchanged.

If the `K_READABLE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts to read access only.

If the `K_FREEZE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts as non-accessible.

If the `K_NOACCESS` flag is set, *xxFlush* writes back all modified data of the target parts and invalidates them.

If the `K_DESTROY` flag is set, *xxFlush* invalidates the target parts without writing back.

If the `K_ASYNC` flag is set, the vm performs the write operations required by the *xxFlush* asynchronously.

If `K_PAGEFAULT` flag is set, *xxFlush* is a result of a page fault. This type of *lcFlush* operation could break the atomicity of a *sgRead/sgWrite* operation running concurrently on the same local cache parts.

The `K_ORDERED` flag can only be used with an *lcFlush* or an *lcSetRights* request. If the flag is set, the *orderNb* argument specifies the *lcFlush* message order number. It allows the kernel to detect the situation when two messages sent by a mapper in one order are received by the kernel in another. For instance, if the mapper grants certain access rights in a *MpGetAccess* reply first, then, processing another request, calls an *lcFlush* to recall the access rights, considering the order numbers, the kernel is aware that the access returned by the *MpGetAccess* reply was already recalled by the mapper, and performs another *MpGetAccess* request.

The *lcSetRights* operation is similar to the *lcFlush* operation, except that it only changes data protections without invalidation and/or writing back: the `K_WRITABLE` flag is ignored, the `K_READABLE` flag sets the target parts read—only, whereas the `K_FREEZE` and `K_NOACCESS` flags set the target parts to non-accessible. The `K_COPYBACK`, `K_ASYNC` and `K_DESTROY` flags are prohibited.

The *sgSyncAll* operation writes back asynchronously all dirty parts of all local caches on the site, except the local caches mapped at least once to a region with the `K_NOSYNC` attribute (see *rgnMap(2SEG)*).

**RETURN VALUE**

If successful `K_OK` is returned, otherwise a negative error code is returned.

**ERRORS**

[`K_EFAULT`] Some of the arguments provided are outside the caller's address space.



[K_EINVAL]	An inconsistent actor capability was provided.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EROUND]	<i>address</i> or <i>size</i> isn't page-aligned.
[K_EROUND]	<i>lcdesc-&gt;startOffset</i> or <i>lcdesc-&gt;size</i> isn't fragment-aligned.
[K_EROUND]	<i>segdesc-&gt;startOffset</i> or <i>segdesc-&gt;size</i> isn't fragment-aligned.
[K_EADDR]	Some or all the addresses from the target address range are invalid.
[K_EOFFSET]	Tried to flush a segment outside the valid offset range in a segment, as returned by <i>vmStat</i> .
[K_EINVAL]	The <i>flags</i> argument contains invalid flag values.
[K_EBUSY]	Tried to invalidate or destroy a no-demand (mapped to a region with K_NODEMAND attribute) physical memory.
[K_EFAIL]	An <i>ipcCall</i> transaction failed during the remote <i>lcFlush</i> or <i>lcSetRights</i> .
[K_EMAPPER]	The mapper doesn't respect the vm/mapper protocol.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site ( *vmFlush* only).

The *sgFlush* and *sgSyncAll* calls remain in the interface for backward compatibility. They will be removed in a future release.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpGetAccess(2SEG)* , *MpPullIn(2SEG)* , *MpPushOut(2SEG)* , *rgnMap(2SEG)* , *lcOpen(2SEG)* , *lcFillZero(2SEG)* , *lcTrunc(2SEG)*

<b>NAME</b>	lcOpen, lcClose, lcCap – Find or create a local cache object for a segment; Release a local cache object; Return the capability of a local cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcOpen(KnCap * sgcap, VmFlags flags, int * cachelip);  int lcCap(int cacheli, KnCap * lccap);  int lcClose(int cacheli);</pre>				
<b>FEATURES</b>	MEM_VM				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcOpen</i> call allows the caller to force the kernel to use the same local cache object for the segment specified until a subsequent <i>lcClose</i> call is received.</p> <p>The <i>sgcap</i> argument is a pointer to the target segment capability. The <i>flags</i> argument must be zero. This argument will be used in future extensions of the interface.</p> <p>The <i>lcOpen</i> call returns into the variable pointed to by the <i>cachelip</i> argument, a local id of the local cache object corresponding to the target segment. The caller can then use the local id in a number of subsequent <i>lcRead</i>, <i>lcWrite</i> (see <i>lcRead(2SEG)</i>) and <i>lcCap</i> (see below) calls. The local id must be released using the <i>lcClose</i> call and should not be used any more.</p> <p>The <i>lcCap</i> call finds the local cache object specified by its local id ( <i>cacheli</i> argument) and returns the local cache capability in the variable pointed to by the <i>lccap</i> argument.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors. An attempt to pass an invalid local id may produce unpredictable system behavior.				
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>The system is out of resources.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_ENOMEM]	The system is out of resources.
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.				
[K_ENOMEM]	The system is out of resources.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>vmStat(2K)</i> , <i>MpGetAccess(2SEG)</i> , <i>MpPullIn(2SEG)</i> , <i>rgnMap(2SEG)</i> , <i>lcFlush(2SEG)</i> , <i>lcFillZero(2SEG)</i> , <i>lcTrunc(2SEG)</i>				

<b>NAME</b>	lcPushData – push data from a source local cache to a target local cache
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; #include &lt;mem/chMapper.h&gt; int lcPushData(KnObjDesc *source, VmFlags sourceFlags, KnObjDesc *target, VmFlags targetFlags, MpPullInId pullInId, unsigned long ordinalNb);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <code>source</code> argument points to a <i>KnObjDesc</i> structure whose members are the following:</p> <pre>KnCap      dataObject ; VmOffset   startOffset ; VmSize     size ;</pre> <p>where the <i>dataObject</i> field is the local cache capability, the <i>startOffset</i> field is the start offset of the part in the local cache, and the <i>size</i> field is the part size. The <i>sourceFlags</i> parameter indicates the new access level for the local cache which is the source cache for the push data operation.</p> <p>If the <code>K_READABLE</code> flag is set, <i>lcPushData</i> sends the requested data to the target local cache and sets the source access to read-only.</p> <p>If the <code>K_NOACCESS</code> flag is set, <i>lcPushData</i> sends the requested data to the target local cache and discards the source data.</p> <p>If the <code>K_PAGEFAULT</code> flag is set, <i>lcPushData</i> is the result of a page fault. This operation could break the atomicity of <i>ansgRead/sgWrite</i> operation running concurrently on the same local cache parts.</p> <p>If the <code>K_ORDERED</code> flag is set, the <i>lcPushData</i> request is an ordered message.</p> <p>The <i>target</i> argument points to a <i>KnObjDesc</i> structure which contains the capability of a local cache which is requesting data for an overlapping area of the same segment. The <i>targetFlags</i> parameter can be used to pass access information to the target site. The <i>pullInId</i> argument is provided in a <i>MpIn</i> request from the target site.</p> <p>The <i>ordinalNb</i> argument specifies the order number of the request. For each local cache the kernel maintains a variable containing the order number expected for the next ordered message. The variable is initialized with zero. When the kernel receives an ordered message for a local cache, it compares the message order number with the order number expected. If the message order number is greater than or equal to the one expected, the kernel sets the expected order number to the message order number plus one and then processes the message. If the</p>

message order number is less than the one expected, the *lcPushData* request is processed, but the expected order number is not modified.

This operation is applied even if the source cache is managed by a remote site. The effect of *lcPushData* on the source cache is similar to *lcFlush* using the same flags. The difference is that *lcFlush* causes modified data to be pushed to the mapper.

**RETURN VALUE** If successful, *lcPushData* returns K\_OK, otherwise a negative error code is returned. The number of bytes pushed to the target cache is returned in *target->size*.

**RESTRICTIONS** The current implementation of *lcPushData* restricts the pushed data to one page.

**ERRORS** [K\_EBUSY] Tried to flush a part of a segment locked in memory.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `lcFlush(2SEG)`, `MpPullIn(2SEG)`

<b>NAME</b>	lcRead, lcWrite – Read data through a local cache into an actor address space; Write data from an actor address space through a local cache
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcRead(int srccheli, VmOffset srcoffseti, KnCap * dstactorcap, VmAdd dstaddress, VmSize * sizep);  int lcWrite(int dstcheli, VmOffset dstoffset, KnCap * srcactorcap, VmAddr srcaddress, VmSize * sizep);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcRead</i> function reads data from a segment through a local cache into the actor address space specified by <i>dstactorcap</i> - a pointer to the actor capability. If <i>dstactorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>dstactorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The source data is specified by the <i>srccheli</i> argument, which is the local id of the local cache corresponding to the source segment and the <i>srcoffset</i> argument, which is the data start offset within the segment. The destination is specified by the <i>dstaddress</i> argument, which is the start address in the destination actor address space.</p> <p>The <i>lcWrite</i> function writes data from an actor address space into a segment through a local cache. The source address space is specified by <i>srcactorcap</i>, which is a pointer to the actor capability. If <i>srcactorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>srcactorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The source data is specified by the <i>srcaddress</i> argument, which is the data start address in the source actor address space. The destination is specified by the <i>dstcheli</i> argument, which is the local id of the local cache corresponding to the destination segment and the <i>dstoffset</i> argument, which is the data start offset within the segment.</p> <p>The data size is specified by the <i>sizep</i> argument - a pointer to the variable containing the size of data.</p> <p>If any error occurs during an <i>lcRead</i> or <i>lcWrite</i> operation, the number of bytes read or written before the error occurred is returned to the variable pointed to by the <i>sizep</i> argument.</p>
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.
<b>ERRORS</b>	[K_EFAULT]                      Some of the arguments provided are outside the caller's address space.

[K_EINVAL]	An inconsistent actor capability was provided.
[K_EUNKNOWN]	<i>dstactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.
[K_EOFFSET]	Attempted to access a segment outside its valid offset range as returned by <i>vmStat</i> .
[K_EFAULT]	Attempted to write to a read-only region.
[K_EFAULT]	The read or write buffer is outside any allocated region.
[K_EMAPPER]	The mapper doesn't respect the vm/mapper protocol.

**RESTRICTIONS**

The caller, the target actor and the local cache must be located on the same site.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`lcOpen(2SEG)` , `lcClose(2K)`

<b>NAME</b>	lcFlush, lcSetRights, sgFlush, sgSyncAll, vmFlush – flush local cache(s)
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcFlush(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int lcSetRights(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int sgFlush(KnObjDesc * segdesc, VmFlags flags);  int sgSyncAll(void);  int vmFlush(KnCap * actorcap, VmAddr address, VmSize size, VmFlags flags);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcFlush</i> call performs a flush operation on a range of a local cache. The <i>lcdesc</i> argument points to a <i>KnObjDesc</i> structure whose members are the following:</p> <pre>KnCap      dataObject ; VmOffset   startOffset ; VmSize     size ;</pre> <p>where the <i>dataObject</i> field is the local cache capability, the <i>startOffset</i> field is the range start offset in the local cache, and the <i>size</i> field is the range size. Both <i>startOffset</i> and <i>size</i> must be fragment-aligned. The size of the fragment is implementation-dependent and is usually equal to the virtual page size divided by 8 (number of bits in one byte).</p> <p>For <i>lcFlush</i>, the target local cache is directly specified by its capability; the flush operation is applied even if the cache is managed by a remote site.</p> <p>The <i>sgFlush</i> call takes the same arguments, except the <i>dataObject</i> field specifies the capability of the cached segment. For <i>sgFlush</i>, the target local cache is indirectly specified by the capability of the corresponding segment; the flush operation implicitly applies to the segment's cache(s) located on the caller's site.</p> <p>The <i>vmFlush</i> call performs a flush operation on the ranges of the local caches mapped to an address range of an actor address space. The address range must be page-aligned.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <i>K_MYACTOR</i>, the address space of the current actor is used. If <i>actorcap</i> is <i>K_SVACTOR</i>, the supervisor address space is used.</p> <p>The <i>flags</i> argument specifies the mode of the flush and the upper access rights for the target parts after the flush operation.</p>

If the K\_COPYBACK flag is set, *xxFlush* writes back any (even clean) cached data of the target parts and keeps the parts access rights unchanged.

If the K\_WRITABLE flag is set, *xxFlush* writes back all modified data of the target parts and keeps the parts access rights unchanged.

If the K\_READABLE flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts to read access only.

If the K\_FREEZE flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts as non-accessible.

If the K\_NOACCESS flag is set, *xxFlush* writes back all modified data of the target parts and invalidates them.

If the K\_DESTROY flag is set, *xxFlush* invalidates the target parts without writing back.

If the K\_ASYNC flag is set, the vm performs the write operations required by the *xxFlush* asynchronously.

If K\_PAGEFAULT flag is set, *xxFlush* is a result of a page fault. This type of *lcFlush* operation could break the atomicity of a *sgRead/sgWrite* operation running concurrently on the same local cache parts.

The K\_ORDERED flag can only be used with an *lcFlush* or an *lcSetRights* request. If the flag is set, the *orderNb* argument specifies the *lcFlush* message order number. It allows the kernel to detect the situation when two messages sent by a mapper in one order are received by the kernel in another. For instance, if the mapper grants certain access rights in a *MpGetAccess* reply first, then, processing another request, calls an *lcFlush* to recall the access rights, considering the order numbers, the kernel is aware that the access returned by the *MpGetAccess* reply was already recalled by the mapper, and performs another *MpGetAccess* request.

The *lcSetRights* operation is similar to the *lcFlush* operation, except that it only changes data protections without invalidation and/or writing back: the K\_WRITABLE flag is ignored, the K\_READABLE flag sets the target parts read—only, whereas the K\_FREEZE and K\_NOACCESS flags set the target parts to non-accessible. The K\_COPYBACK, K\_ASYNC and K\_DESTROY flags are prohibited.

The *sgSyncAll* operation writes back asynchronously all dirty parts of all local caches on the site, except the local caches mapped at least once to a region with the K\_NOSYNC attribute (see *rgnMap(2SEG)*).

**RETURN VALUE**

If successful K\_OK is returned, otherwise a negative error code is returned.

**ERRORS**

[K\_EFAULT]                      Some of the arguments provided are outside the caller's address space.



[K_EINVAL]	An inconsistent actor capability was provided.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EROUND]	<i>address</i> or <i>size</i> isn't page-aligned.
[K_EROUND]	<i>lcdesc-&gt;startOffset</i> or <i>lcdesc-&gt;size</i> isn't fragment-aligned.
[K_EROUND]	<i>segdesc-&gt;startOffset</i> or <i>segdesc-&gt;size</i> isn't fragment-aligned.
[K_EADDR]	Some or all the addresses from the target address range are invalid.
[K_EOFFSET]	Tried to flush a segment outside the valid offset range in a segment, as returned by <i>vmStat</i> .
[K_EINVAL]	The <i>flags</i> argument contains invalid flag values.
[K_EBUSY]	Tried to invalidate or destroy a no-demand (mapped to a region with K_NODEMAND attribute) physical memory.
[K_EFAIL]	An <i>ipcCall</i> transaction failed during the remote <i>lcFlush</i> or <i>lcSetRights</i> .
[K_EMAPPER]	The mapper doesn't respect the vm/mapper protocol.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site ( *vmFlush* only).

The *sgFlush* and *sgSyncAll* calls remain in the interface for backward compatibility. They will be removed in a future release.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpGetAccess(2SEG)* , *MpPullIn(2SEG)* , *MpPushOut(2SEG)* , *rgnMap(2SEG)* , *lcOpen(2SEG)* , *lcFillZero(2SEG)* , *lcTrunc(2SEG)*

<b>NAME</b>	lcStat, sgStat – get the statistics of a local cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; #include &lt;mem/chMapper.h&gt; int lcStat(KnCap * lccap, KnLcStat * stat);  int sgStat(KnCap * segcap, KnLcStat * stat);</pre>				
<b>FEATURES</b>	MEM_VM				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcStat</i> and <i>sgStat</i> functions get the statistics of the local cache specified either by <i>lccap</i> - a local cache capability, or by <i>segcap</i> - a segment capability. In the latter case, the statistics of the cache of the segment on the current site are returned.</p> <p>The <i>KnLcStat</i> structure describes the statistics associated with a local cache, as follows:</p> <pre>VmSize  physMem ; VmSize  lockMem ; KnCap   segcap  ; KnCap   lccap   ;</pre> <p>The <i>physMem</i> field specifies the physical memory size currently allocated for the local cache.</p> <p>The <i>lockMem</i> field specifies the physical memory size currently fixed for the local cache.</p> <p>The <i>segcap</i> field specifies the capability of the corresponding segment. It is returned by the <i>lcStat</i> call only when the caller is a system actor or when the current thread executes in privileged mode.</p> <p>The <i>lccap</i> field specifies the capability of the segment's local cache on the current site. It is returned by the <i>sgStat</i> call only when the caller is a system actor or when the current thread executes in privileged mode.</p>				
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EFAULT]                   Some of the arguments provided are outside the caller's address space.</p> <p>[K_EUNDEF]                   The segment specified is not cached on the site.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				

**SEE ALSO**

lcFlush(2SEG)

<b>NAME</b>	lcTrunc – shape the end of a local cache						
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcTrunc(KnObjDesc *lcdesc, VmOffset zerooffset, VmFlags flags, unsigned long ordinalnb);</pre>						
<b>FEATURES</b>	MEM_VM						
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcTrunc</i> function grants write access to the [<i>lcdesc-&gt;startOffset</i>, <i>lcdesc-&gt;startOffset</i> + <i>lcdesc-&gt;size</i>] range and recalls all access rights from the [<i>lcdesc-&gt;startOffset</i> + <i>lcdesc-&gt;size</i>, <i>K_MAXVMOFFSET</i>] range of the local cache specified by the <i>lcdesc-&gt;dataObject</i> capability.</p> <p>Then, <i>lcTrunc</i> caches the [<i>lcdesc-&gt;startOffset</i>, <i>zerooffset</i>) range in the local cache and zeros the [<i>zerooffset</i>, <i>lcdesc-&gt;startOffset</i> + <i>lcdesc-&gt;size</i>) range (if not empty, in other words, <i>zerooffset</i> == <i>lcdesc-&gt;startOffset</i> + <i>lcdesc-&gt;size</i>).</p> <p>Finally, <i>lcTrunc</i> marks the [<i>lcdesc-&gt;startOffset</i>, <i>lcdesc-&gt;startOffset</i> + <i>lcdesc-&gt;size</i>) range as modified.</p> <p>If <i>lcdesc-&gt;size</i> is equal to zero the <i>lcTrunc</i> call simply invalidates the [<i>lcdesc-&gt;startOffset</i>, <i>K_MAXVMOFFSET</i>] range.</p> <p>Both <i>lcdesc-&gt;startOffset</i> and <i>lcdesc-&gt;size</i> must be fragment-aligned. The size of the fragment is implementation-dependent and usually equal to the virtual page size divided by 8 (number of bits in one byte).</p> <p>If the <i>K_ORDERED</i> flag is set in the <i>flags</i> argument, the <i>ordinalNumber</i> argument specifies the <i>lcTrunc</i> message order number. It allows the kernel to detect the situation when two messages sent by a mapper in one order are received by the kernel in another. For instance, if the mapper grants certain access rights in a <i>MpGetAccess</i> reply first, then, processing another request, calls an <i>lcTrunc</i> to recall the access rights, using the order numbers, the kernel is aware that the access returned by the <i>MpGetAccess</i> reply was already recalled by the mapper, and performs another <i>MpGetAccess</i> request.</p> <p>The <i>lcTrunc</i> operation cannot be blocked by either a pure <i>MpGetAccess</i> or by an impure (combined with a <i>MpGetAccess</i>) <i>MpPullIn</i> up-call.</p>						
<b>RETURN VALUE</b>	If successful <i>K_OK</i> is returned, otherwise a negative error code is returned.						
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td>An inconsistent local cache capability was provided.</td> </tr> <tr> <td style="vertical-align: top;">[K_EUNKNOWN]</td> <td><i>lcdesc-&gt;dataObject</i> does not specify a reachable local cache.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_EINVAL]	An inconsistent local cache capability was provided.	[K_EUNKNOWN]	<i>lcdesc-&gt;dataObject</i> does not specify a reachable local cache.
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.						
[K_EINVAL]	An inconsistent local cache capability was provided.						
[K_EUNKNOWN]	<i>lcdesc-&gt;dataObject</i> does not specify a reachable local cache.						

[K_EROUND]	<i>lcdesc-&gt;startOffset</i> or <i>lcdesc-&gt;size</i> isn't fragment-aligned.
[K_EOFFSET]	Tried to truncate a segment outside the valid offset range in a segment, as returned by <i>vmStat</i> .
[K_EOFFSET]	<i>zerooffset</i> is out of [ <i>lcdesc-&gt;startOffset</i> , <i>lcdesc-&gt;startOffset</i> + <i>lcdesc-&gt;size</i> ] range.
[K_EINVAL]	The <i>flags</i> argument contains invalid flag values.
[K_EBUSY]	Tried to invalidate or destroy a no-demand (mapped to a region with the K_NODEMAND attribute) physical memory.
[K_EFAIL]	An <i>ipcCall</i> transaction failed during the remote <i>lcTrunc</i> .
[K_EMAPPER]	The mapper doesn't respect the vm/mapper protocol.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpGetAccess(2SEG)*, *MpPullIn(2SEG)*, *rgnMap(2SEG)*, *lcOpen(2SEG)*, *lcFlush(2SEG)*, *lcFillZero(2SEG)*

<b>NAME</b>	lcRead, lcWrite – Read data through a local cache into an actor address space; Write data from an actor address space through a local cache
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcRead(int srccheli, VmOffset srcoffseti, KnCap * dstactorcap, VmAddr dstaddress, VmSize * sizep);  int lcWrite(int dstcheli, VmOffset dstoffset, KnCap * srcactorcap, VmAddr srcaddress, VmSize * sizep);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcRead</i> function reads data from a segment through a local cache into the actor address space specified by <i>dstactorcap</i> - a pointer to the actor capability. If <i>dstactorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>dstactorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The source data is specified by the <i>srccheli</i> argument, which is the local id of the local cache corresponding to the source segment and the <i>srcoffset</i> argument, which is the data start offset within the segment. The destination is specified by the <i>dstaddress</i> argument, which is the start address in the destination actor address space.</p> <p>The <i>lcWrite</i> function writes data from an actor address space into a segment through a local cache. The source address space is specified by <i>srcactorcap</i>, which is a pointer to the actor capability. If <i>srcactorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>srcactorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The source data is specified by the <i>srcaddress</i> argument, which is the data start address in the source actor address space. The destination is specified by the <i>dstcheli</i> argument, which is the local id of the local cache corresponding to the destination segment and the <i>dstoffset</i> argument, which is the data start offset within the segment.</p> <p>The data size is specified by the <i>sizep</i> argument - a pointer to the variable containing the size of data.</p> <p>If any error occurs during an <i>lcRead</i> or <i>lcWrite</i> operation, the number of bytes read or written before the error occurred is returned to the variable pointed to by the <i>sizep</i> argument.</p>
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.
<b>ERRORS</b>	[K_EFAULT]                      Some of the arguments provided are outside the caller's address space.

[K_EINVAL]	An inconsistent actor capability was provided.
[K_EUNKNOWN]	<i>dstactorcap</i> or <i>srcactorcap</i> does not specify a reachable actor.
[K_EOFFSET]	Attempted to access a segment outside its valid offset range as returned by <i>vmStat</i> .
[K_EFAULT]	Attempted to write to a read-only region.
[K_EFAULT]	The read or write buffer is outside any allocated region.
[K_EMAPPER]	The mapper doesn't respect the vm/mapper protocol.

**RESTRICTIONS**

The caller, the target actor and the local cache must be located on the same site.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`lcOpen(2SEG)` , `lcClose(2K)`

**NAME** MpCreate – create a temporary segment at the default mapper

**SYNOPSIS**

```
#include <mem/chMem.h>
#include <mem/chMapper.h>
MpCreate
request annex(KnMpCreate structure):
    int         service;
    VmFlags     options;
response annex (KnMpCreateReply structure):
    int         diag;
    KnCap       segcap;
    VmFlags     options;
```

**FEATURES** MEM\_VM

**DESCRIPTION** **Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.

In order to create a temporary segment at the default mapper, the kernel memory management system performs an *MpCreate* message transaction. In other words, it sends, using *ipcCall(2K)*, an *MpCreate* request message to a port or group of ports identifying the mapper.

The request message consists of an annex (no body) whose head matches the *KnMpCreate* structure defined above. The *service* field of this structure must be set to KN\_MPCREATE.

The mapper replies with a message, also consisting of an annex only, whose head must match the *KnMpCreateReply* structure. The *diag* field is the operation return code. The *diag* must be either 0 (K\_OK) or a negative number. If a negative error code is returned, the kernel returns it to the original kernel call (if any).

The *segcap* field contains the segment capability created.

The *options* field of the return message must be zero for backward compatibility with future versions.

**ATTRIBUTES** See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *MpGetAccess(2SEG)*, *MpPullIn(2SEG)*, *MpPushOut(2SEG)*, *MpRelease(2SEG)*



<b>NAME</b>	MpGetAccess – get access to data through a local cache
<b>SYNOPSIS</b>	<pre> #include &lt;mem/chMem.h&gt; #include &lt;mem/chMapper.h&gt; MpGetAccess request annex (KnMpGetAccess structure):     int          service ;     KnKey        segkey ;     KnCap        lccap ;     VmOffset     accessOffset ;     VmSize       accessSize ;     VmFlags      requiredAccess ;     VmSize       requiredAccessOffset ;     VmSize       requiredAccessSize ; response annex (KnMpGetAccessReply structure):     int          diag ;     VmFlags      grantedAccess ;     unsigned long ordinalNumber ;     VmOffset     returnAccessOffset ;     VmSize       returnAccessSize ;     VmSize       inClusterSize ;     VmSize       outClusterSize ; </pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>In order to demand access rights for a range of a segment through a local cache, the kernel memory management system performs an <i>MpGetAccess</i> message transaction. In other words it sends, using <i>ipcCall(2K)</i>, an <i>MpGetAccess</i> request message to the mapper managing the segment. The mapper is identified by the <i>ui</i> field of the segment capability.</p> <p>The request message consists of an annex (no body) whose head matches the <i>KnMpGetAccess</i> structure defined above. The <i>service</i> field of this structure must be set to KN_MPGETACCESS.</p> <p>The kernel demands access rights to <i>accessSize</i> bytes of data using the <i>accessOffset</i>, the starting offset in the segment is specified by <i>segkey</i> on the invoked mapper,</p>

for the local cache specified by *lccap*. Both *accessOffset* and *accessSize* are always page-aligned.

The kernel also specifies in the message that access rights for *requiredAccessSize* bytes of data with the *requiredAccessOffset* starting offset is mandatory.

The kernel guarantees that  $[accessOffset, accessOffset + accessSize]$  includes  $[requiredAccessOffset, requiredAccessOffset + requiredAccessSize]$ .

The kernel can send a get access request when it performs one of the following operations:

- *lc(sg)Read* or a *lc(sg)Write* kernel calls. In this case  $[requiredAccessOffset, requiredAccessOffset + requiredAccessSize]$  specifies the exact (in bytes) boundaries of the operation.

The kernel can perform a *lc(sg)Read/Write* call without a *MpGetAccess* transaction if it already has access to the required range.

- any other operation which acts on a range of the segment through a virtual address space as a page fault or *vmLock*, *vmCopy*, *rgnSetPaging*, *rgnInit*, for example.

In this, and only this, case the *K\_PAGEFAULT* flag is set in the *requiredAccess* field.

The kernel can ask for access rights to the operation range (the operation range of a page fault is the whole faulty page) using a number of consecutive *MpGetAccess* transactions. The *requiredAccessOffset* of each transaction is page-aligned and  $requiredSize \% vmPageSize()$  is equal to 1 (the last byte of the required range is the first byte of a page).

The kernel can perform the operation without an *MpGetAccess* request for a particular page overlapping the operation range if it has already the requested access to the whole page.

If the *K\_WRITABLE* flag in the *requiredAccess* field is set, write access to the data is requested, otherwise read-only access is requested.

The mapper must reply with a message, consisting of an annex (no body) whose head matches the *KnMpGetAccessReply* structure.

The *diag* field is the transaction return code. The *diag* field must be either 0 (*K\_OK*), or a negative number. When a negative error code is returned, there are two possibilities:

- *returnAccessSize* is equal to zero. In this case the kernel immediately gives up the processing of the original operation and returns from the original kernel operation with a *diag* error code.

- *returnAccessSize* is positive. In this case the kernel reduces the original operation range until *returnAccessOffset + returnAccessSize*, completes the operation and returns from the original kernel operation with a *diag* error code. Also, if it's allowed by the interface of the operation (*lcRead* or *lcWrite* calls), the number of bytes effectively processed will be returned.

If the original kernel operation is a page fault an error provokes exception processing at the end of the page fault handler execution (see *svExcHandler(2K)*).

The *K\_MPIVERS* bitmask defines a bit-field of the *grantedAccess* field which specifies the version of the *MpGetAccess* protocol supported by the mapper. The protocol version has to be equal to *K\_MPIVER1*.

The mapper grants access to *returnAccessSize* bytes of data with the starting offset in the segment equal to *returnAccessOffset*. The mapper has to guarantee that *[returnAccessOffset, returnAccessOffset + returnAccessSize]* is included in *[accessOffset, accessOffset + accessSize]* but includes *[requiredAccessOffset, requiredAccessOffset + requiredAccessSize]*. Note that if *diag* is negative (see above) the last condition can be not respected.

The kernel requires *returnAccessOffset* and *returnAccessSize* to be "fragment"-aligned. The size of the fragment is implementation—dependent but usually equal to the virtual page size divided by 8 (number of bits in one byte).

The alignment requirements above mean that the mapper can return access rights to a partial page. The kernel implements the following policy with respect to pages with partial access rights:

- The kernel maps the page to the target virtual address if required by the target operation. The kernel page fault handler can map a page with partial access rights to the faulty virtual address after the corresponding *MpGetAccess* return. Note that in the case of another page fault on a page with partial access rights, the kernel will call *MpGetAccess* again.
- The kernel doesn't guarantee that the page compliment (the ranges of the page without access rights) will be included in a subsequent *mpPushOut* request even if it was modified using mmu. Neither does the kernel guarantee that the page complement will not be included in any subsequent *mpPushOut* request.

If the mapper grants write access, it sets the *K\_WRITABLE* flag in the *grantedAccess* field, otherwise read only access is granted. The mapper could grant write access in response to a read-only access request, but it must grant write access in response to a write access request.

If the `K_ORDERED` flag is set in the *grantedAccess* field, the *ordinalNumber* field specifies the reply message order number. It allows the kernel to detect the situation when two messages sent by a mapper in one order are received by the kernel in another. For instance, suppose the mapper grants access rights in an *MpGetAccess* reply first, and then, processing another request, calls an *lcFlush* to recall the access rights. In this case, the order numbers indicate to the kernel that the access returned by the *MpGetAccess* reply was already recalled by the mapper, and performs a *MpGetAccess* request again.

If the `K_STALE` flag is set in the *grantedAccess* field, any data already cached from the returned range are considered as stale and will be destroyed by the kernel upon receipt of the reply message.

The mapper may also set the `K_RELEASE` flag (in the *grantedAccess* field). In this case the memory management has to perform an *MpRelease* request prior to destroying the local cache.

If the `K_GETATTR` flag is set in the *requiredAccess* field of the request message, the current *MpGetAccess* request is combined with a "get segment attributes" request. When the mapper replies to a combined request it may return the segment clustering attributes. If the mapper sets the `K_CLUSTER` flag (in the *grantedAccess* field) the *inClusterSize* and *outClusterSize* fields define the segment clustering attributes. The *inClusterSize* defines the optimal (from the mapper point of view) data size for any future *MpPullIn* operation performed on the segment. The *outClusterSize* defines the optimal data size for any future *MpPushOut* operation. Both *inClusterSize* and *outClusterSize* have to be a power of two and page-aligned.

If the kernel detects that the mapper is not adhering to the protocol described above, the `K_EMAPPER` error is returned to the original kernel operation. If the original kernel operation is a page fault an error provokes exception processing at the end of the page fault handler execution (see *svExcHandler(2K)*).

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpPullIn(2SEG)*, *MpPushOut(2SEG)*, *rgnMap(2SEG)*, *lcOpen(2SEG)*, *lcFlush(2SEG)*, *lcFillZero(2SEG)*, *lcTrunc(2SEG)*

<b>NAME</b>	MpPullIn – read data into a local cache
<b>SYNOPSIS</b>	<pre> #include &lt;mem/chMem.h&gt; #include &lt;mem/chMapper.h&gt; MpPullIn request annex (KnMpInle structure):     int          service ;     KnCap        segkey ;     KnCap        lccap ;     VmOffset     accessOffset ;     VmSize       accessSize ;     VmFlags      requiredAccess ;     VmOffset     requiredAccessOffset ;     VmSize       requiredAccessSize ;     VmOffset     dataOffset ;     VmSize       dataSize ;     VmOffset     requiredDataOffset ;     VmSize       requiredDataSize ;     MpPullInId  pullInId ; response annex (KnMpInReply structure) :     int          diag ;     VmFlags      grantedAccess ;     unsigned     ordinalNumber ;     VmOffset     returnAccessOffset ;     VmSize       returnAccessSize ;     VmSize       inClusterSize ;     VmSize       outClusterSize ;     VmOffset     returnDataOffset ;     VmSize       returnDataSize ; response body:     unsigned char data [bodysize] ; </pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.

In order to demand the contents of a range of a local cache to be updated from the corresponding segment, the kernel memory management system performs a *MpPullIn* message transaction. It sends, using *ipcCall(2K)* a *MpPullIn* request message to the mapper managing the segment. The mapper is identified by the *ui* field of the segment capability.

The request message consists of an annex (no body) whose head matches the *KnMpIn* structure defined above. The *service* field of this structure is set to *KN\_MPPULLIN*.

The kernel demands *dataSize* bytes of data to be updated with the starting offset specified by *dataOffset*, from the segment specified by *segkey* on the invoked mapper, in the local cache specified by *lccap*. The kernel also specifies in the message that the update of *requiredDataSize* bytes of data with the *requiredDataOffset* starting offset is mandatory. The kernel guarantees that *[dataOffset, dataOffset + dataSize]* includes *[requiredDataOffset, requiredDataOffset + requiredDataSize]*. The *dataOffset*, *dataSize*, *requiredDataOffset* and *requiredDataSize* parameters are always page aligned. The *pullInId* identifier is the *MpPullIn* message transaction identifier used in *lcPushData(2SEG)*.

The mapper must reply (using *ipcReply(2K)*) with a message whose annex head matches the *KnMpInReply* structure, and whose body (if any) contains the data to be put in the local cache.

The *diag* field is the operation return code. The *diag* field must be either 0 (*K\_OK*), or a negative number. If a negative error code is returned, the kernel returns it to the original kernel call.

The *returnDataSize* field specifies the size of the data to be put in the local cache, and the *returnDataOffset* field specifies the starting offset of the data within the segment. The mapper has to guarantee that *[returnDataOffset, returnDataOffset + returnDataSize]* is included in *[dataOffset, dataOffset + dataSize]* but includes *[requiredDataOffset, requiredDataOffset + requiredDataSize]*. Both *returnDataOffset* and *returnDataSize* must be page aligned.

The size of the body may be less or equal to the *returnDataSize* value. If the body size is less than *returnDataSize*, the memory management system will fill the remaining data with default values: zero or scratch. The default value is zero if the mapper sets the *K\_FILLZERO* flag in the *grantedAccess* field.

If the *K\_DIRTY* flag is set in the *grantedAccess* field, the kernel considers the returned data as dirty: the kernel will later perform an *MpPushOut* operation for the data even if the data is not modified on the site.

If *accessSize* is not zero, the *MpPullIn* data request is combined with a *getAccess* access request. The mapper must reply to the access request as described in the *MpGetAccess (2SEG)* manual.

In the case of a combined request, the mapper can defer data return until a subsequent non-combined (pure) *MpPullIn* request. To reply without data, the mapper sets the *returnDataSize* field and the reply body size to zero.

If the mapper replies to a combined *MpPullIn* request with data, it can't use the *K\_DIRTY* and *K\_ORDERED* flags together.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpGetAccess(2SEG)*, *MpPushOut(2SEG)*, *rgnMap(2SEG)*, *lcOpen(2SEG)*, *lcPushData(2SEG)*

<b>NAME</b>	MpPushOut – write data back to mapper
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; #include &lt;mem/chMapper.h&gt; MpPushOut request annex (KnMpOut structure) :     int            service     KnKey          segkey     KnCap          lccap     VmOffset       dataOffset     VmSize         dataSize     VmFlags        options     KnAsynIoDesc  asynIoDesc request body:     unsigned char data [bodysize] response annex (KnMpOutReply structure) :     int            diag</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It <b>MUST NOT</b> be used by any application.</p> <p>In order to write back the contents of a range of a local cache to the corresponding segment, the kernel memory management system performs a <i>MpPushOut</i> message transaction. It sends, using <i>ipcCall(2K)</i> an <i>MpPushOut</i> request message to the mapper managing the segment. The mapper is identified by the <i>ui</i> field of the segment capability.</p> <p>The request message consists of an annex whose head matches the <i>KnMpOut</i> structure defined above, and a body containing the data to be written. The <i>service</i> field is set to KN_MPPUSHOUT.</p> <p>Upon receipt of such a request, the mapper writes data back to the segment specified by <i>segkey</i> from the local cache specified by <i>lccap</i>. If the segment was created using <i>MpCreate(2SEG)</i>, the <i>lccap</i> argument is undefined.</p> <p>The <i>dataOffset</i> field specifies the starting offset of the data in the segment, and the <i>dataSize</i> field the data size. Currently, the <i>dataSize</i> field is always equal to the message body size and "fragment"-aligned. The size of the fragment is implementation—dependent and usually equal to the virtual page size divided by 8 (number of bits in one byte). The <i>dataOffset</i> field is always page-aligned.</p>



If the `K_ASYNC_REQUEST` flag is set in the *options* field, the mapper can perform an asynchronous write (see below). In this case the kernel fills the *asynIoDesc* structure which has to be passed as a parameter for a subsequent *ioDone* call. The *asynIoDesc* structure is opaque to the mapper.

The mapper must reply to a message whose annex head matches the *KnMpOutReply* structure, where the *diag* field is the operation return code. The *diag* field must be 0 (`K_OK`) or 1 (`K_ASYNC_REPLY`) or a negative number. If a negative error code is returned, the kernel returns it to the original kernel call (if any). The `K_ASYNC_REPLY` can be returned by the mapper if, and only if, the `K_ASYNC_REQUEST` flag was set by the kernel.

If `K_ASYNC_REPLY` is returned (the mapper performs an asynchronous write), the mapper must invoke an *ioDone* function in order to notify that a write operation has been performed, as follows:

```
void ioDone (diag, asynIoDesc)
            int      diag ;
            KnAsynIoDesc* asynIoDesc ;
```

The *diag* parameter must be either 0 (`K_OK`) or a negative number. If a negative error code is returned, the kernel returns it to the original kernel call (if any).

The *asynIoDesc* parameter points to the *KnAsynIoDesc* structure corresponding to the *MpPushOut* request.

The *ioDone* function can be called asynchronously (in respect of *ipcCall* reply) from base or interrupt level.

#### RESTRICTIONS

The current kernel implementation never performs asynchronous *MpPushOut* requests to a remote mapper.

The *ioDone* function implementation is only for trusted supervisor actors: an attempt to pass an invalid *asynIoDesc* may produce undefined system behavior.

#### ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

#### SEE ALSO

`MpPullIn(2SEG)`, `MpCreate(2SEG)`, `lcFlush(2SEG)`, `rgnMap(2SEG)`, `lcOpen(2SEG)`

<b>NAME</b>	MpRelease – release a temporary segment or notify a local cache destruction				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; #include &lt;mem/chMapper.h&gt; MpRelease request annex (KnMpRelease structure) :     int     service ;     KnKey   segkey ;     KnCap   lccap ; response annex (KnMpReleaseReply structure) :     int     diag ;</pre>				
<b>FEATURES</b>	MEM_VM				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>When a temporary segment previously created using <i>MpCreate(2SEG)</i> is no longer in use, an <i>MpRelease</i> message transaction is performed by the kernel. The kernel also performs this request, when destroying a local cache for which the K_RELEASE flag was set by the mapper (see <i>MpGetAccess(2SEG)</i>). This request message is sent to the mapper managing the segment, using <i>ipcCall(2K)</i>. The mapper is identified by the <i>ui</i> field of the segment capability.</p> <p>The request message consists of an annex (no body) whose head matches the <i>KnMpRelease</i> structure defined above.</p> <p>The <i>service</i> field is set to KN_MPRELEASE.</p> <p>The <i>segkey</i> field identifies the segment and the <i>lccap</i> field, if needed, identifies the local cache just destroyed.</p> <p>The mapper must reply with a message whose annex head matches the <i>KnMpReleaseReply</i> structure, where the <i>diag</i> field is the operation return code. The <i>diag</i> must be either 0 (K_OK) or a negative number. If a negative error code is returned, the kernel returns it to the original kernel call (if any).</p>				
<b>ATTRIBUTES</b>	<p>See <i>attributes(5)</i> for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>MpCreate(2SEG)</i> , <i>MpGetAccess(2SEG)</i>				

<b>NAME</b>	pageIoDone – Inform nucleus I/O is complete on a page list				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; void pageIoDone(KnPage *page, int diag);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>pageIoDone</i> function must be invoked by the Proxy-Mapper when a write operation previously triggered by a <i>pushOutAsyn</i> up-call is done, or when a read operation following a <i>dcGetPages</i> is done. The <i>page</i> argument specifies a page list which may be a sub-list of the original lists passed to <i>pushOutAsyn</i> or returned to <i>dcGetPages</i>. The Proxy-Mapper may split the original lists into a number of sub-lists and perform the <i>pageIoDone</i> call for each sub-list independently. The <i>diag</i> argument specifies the result of the operation. If successful it must be set to K_OK, otherwise it must be an error code.</p> <p>Programmers should take care to invoke the <i>pageUnmap</i> nucleus call before invoking the <i>pageIoDone</i> call. Doing it the other way round is likely to lead to unpredictable behaviour of the system.</p> <p>Sub-lists passed to <i>pageUnmap</i> and to <i>pageIoDone</i> do not need to be identical.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	NONE				
<b>ERRORS</b>	No error messages are returned.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>pageMap(2SEG)</i> , <i>PxmPushOutAsyn(2SEG)</i> , <i>dcGetPages(2SEG)</i>				

<b>NAME</b>	pageMap, pageUnmap – Map a list of pages in the current actor address space; Unmap a list of pages
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int pageMap(KnPage * page, VmFlags flags, VmAddr * addr, VmSize * size, VmOffset * offset);  void pageUnmap(KnPage * page);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>pageMap</i> call contiguously maps the page list specified by the <i>page</i> argument in the supervisor address space The <i>flags</i> argument consists of two parts:</p> <ul style="list-style-type: none"> <li>■ The K_KNPAGE_TYPE bit field which specifies the type of page list: <ul style="list-style-type: none"> <li>K_KNPAGE_IN                      For <i>getAcc</i> , <i>pullIn</i> or <i>dcGetPages</i> page lists</li> <li>K_KNPAGE_OUT                    For <i>pushOutAsyn</i> page lists</li> </ul> </li> <li>■ The K_NOWAITFORMEMORY flag; if this flag is set, the call will return immediately with the <i>K_ENOMEM</i> error if there is no virtual address space available. If not set, it will wait for virtual memory to become available. This flag must be set if the <i>pageMap</i> call is invoked from interrupt level.</li> </ul> <p>The <i>pageMap</i> call sets the <i>addr</i> output argument to the virtual address to which the page list is mapped. It also sets the <i>size</i> output argument to the size of the data, and the <i>offset</i> output argument to the offset of the first page of the list. In the case of a <i>pullIn</i> up-call, or of a <i>dcGetPages</i>, the output argument <i>size</i> is set to the number of pages in the list multiplied by the page size. In the case of a <i>pushOutAsyn</i> up-call, <i>pageMap</i> gives a valid data size which is less than or equal to the number of pages of the list multiplied by the page size (the last page of the list may be partially valid).</p> <p>The <i>pageUnmap</i> nucleus call unmaps a list of pages specified by the <i>page</i> argument which were previously mapped through the <i>pageMap</i> call.</p> <p>It should be noted that the list of pages may be freely split and joined, either before invoking <i>pageMap</i> or <i>pageUnmap</i> . The lists provided are always sorted incrementally by offset.</p>
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.
<b>RETURN VALUE</b>	If successful, K_OK is returned, otherwise an error code is returned.
<b>ERRORS</b>	[K_ENOMEM]                      The system is out of resources.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`PxmGetAcc(2SEG)` , `dcGetPages(2SEG)` , `pageIoDone(2SEG)`

**NAME** | pagePhysAddr – Get the physical address of a page

**SYNOPSIS** | #include <mem/chMem.h>  
 PhAddr pagePhysAddr(KnPage \*page, VmFlags flags, VmSize \*size, VmOffset \*offset);

**FEATURES** | PXM\_EXT

**DESCRIPTION** | **Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.  
 The *pagePhysAddr* call returns the physical address of the physical page specified by the *page* argument. The *flags* argument specifies the type of page:  
 K\_KNPAGE\_IN | Page is used for a *getAcc*, *pullIn* or *dcGetPages* call.  
 K\_KNPAGE\_OUT | Page is used for a *pushOutAsyn* call.  
 When *flags* is set to K\_KNPAGE\_IN, *pagePhysAddr* sets the *size* output argument to the page size. When *flags* is set to K\_KNPAGE\_OUT, *pagePhysAddr* gives the valid data size which is less than or equal to the page size. The *pagePhysAddr* call sets the *offset* output argument to the page offset.

**RESTRICTIONS** | The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE** | This call returns the physical address of the page.

**ERRORS** | none

**ATTRIBUTES** | See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** | *pageMap(2SEG)*

<b>NAME</b>	pageSetDirty – Mark a page as dirty				
<b>SYNOPSIS</b>	#include <mem/chMem.h> void <b>pageSetDirty</b> (KnPage *page);				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.  The <i>pageSetDirty</i> call enables the proxy-mapper to mark as dirty a page being read as part of a <i>pullIn</i> , <i>getAcc</i> up-call or as part of read-ahead processing following a call to <i>dcGetPages</i> . This is more flexible than marking a full page list dirty, as a result of a <i>getAcc</i> or <i>pullIn</i> up-call. Moreover, it also allows the proxy-mapper to mark a page as being dirty during a read-ahead operation following a <i>dcGetPages</i> invocation.				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	NONE				
<b>ERRORS</b>	No error messages are returned.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>pageMap(2SEG)</i> , <i>PxmPushOutAsyn(2SEG)</i> , <i>dcGetPages(2SEG)</i>				

<b>NAME</b>	pageSgId – Get the segment identifier associated with a page				
<b>SYNOPSIS</b>	#include <mem/chMem.h> int <b>pageSgId</b> (KnPage *page);				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It <b>MUST NOT</b> be used by any application.</p> <p>The <i>pageSgId</i> call returns the segment identifier corresponding to the data cache to which the page (or page list) specified by the <i>page</i> argument is attached. This segment identifier is the one defined by the proxy-mapper at <i>dcAlloc(2SEG)</i> invocation.</p> <p>Note that if the page list used is the one provided by a <i>getAcc</i> up-call, <i>pageSgId</i> may return NULL, because conditionally attached pages may have been detached by <i>dcFillZero</i> or <i>dcTrunc</i>.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	This call returns the segment identifier, or NULL if there are no pages attached to the data cache.				
<b>ERRORS</b>	No error messages are returned.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>dcAlloc(2SEG)</i> , <i>dcFlush(2SEG)</i> , <i>dcFillZero(2SEG)</i> , <i>dcTrunc(2SEG)</i> , <i>PxmGetAcc(2SEG)</i> , <i>pageMap(2SEG)</i>				



<b>NAME</b>	pageMap, pageUnmap – Map a list of pages in the current actor address space; Unmap a list of pages
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int pageMap(KnPage * page, VmFlags flags, VmAddr * addr, VmSize * size, VmOffset * offset);  void pageUnmap(KnPage * page);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>pageMap</i> call contiguously maps the page list specified by the <i>page</i> argument in the supervisor address space The <i>flags</i> argument consists of two parts:</p> <ul style="list-style-type: none"> <li>■ The K_KNPAGE_TYPE bit field which specifies the type of page list: <ul style="list-style-type: none"> <li>K_KNPAGE_IN                      For <i>getAcc</i> , <i>pullIn</i> or <i>dcGetPages</i> page lists</li> <li>K_KNPAGE_OUT                    For <i>pushOutAsyn</i> page lists</li> </ul> </li> <li>■ The K_NOWAITFORMEMORY flag; if this flag is set, the call will return immediately with the K_ENOMEM error if there is no virtual address space available. If not set, it will wait for virtual memory to become available. This flag must be set if the <i>pageMap</i> call is invoked from interrupt level.</li> </ul> <p>The <i>pageMap</i> call sets the <i>addr</i> output argument to the virtual address to which the page list is mapped. It also sets the <i>size</i> output argument to the size of the data, and the <i>offset</i> output argument to the offset of the first page of the list. In the case of a <i>pullIn</i> up-call, or of a <i>dcGetPages</i>, the output argument <i>size</i> is set to the number of pages in the list multiplied by the page size. In the case of a <i>pushOutAsyn</i> up-call, <i>pageMap</i> gives a valid data size which is less than or equal to the number of pages of the list multiplied by the page size (the last page of the list may be partially valid).</p> <p>The <i>pageUnmap</i> nucleus call unmaps a list of pages specified by the <i>page</i> argument which were previously mapped through the <i>pageMap</i> call.</p> <p>It should be noted that the list of pages may be freely split and joined, either before invoking <i>pageMap</i> or <i>pageUnmap</i> . The lists provided are always sorted incrementally by offset.</p>
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.
<b>RETURN VALUE</b>	If successful, K_OK is returned, otherwise an error code is returned.
<b>ERRORS</b>	[K_ENOMEM]                      The system is out of resources.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`PxmGetAcc(2SEG)` , `dcGetPages(2SEG)` , `pageIoDone(2SEG)`

<b>NAME</b>	PxmOpen, PxmClose – Open a Data Segment; Close a Data Segment
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; void PxmOpen(KnPxmOpenArgs * openArg, void * cookie);  void PxmClose(KnPxmCloseArgs * closeArgvoid, void * cookie);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>PxmOpen/PxmClose</i> routines are provided by the Proxy-Mapper and are invoked as up-calls by the CHORUS VM. They are invoked as <i>laps</i>; their second argument is the lap cookie as defined at <i>svLapCreate</i> time. The <i>PxmOpen/PxmClose</i> up-calls are invoked by the VM in order to inform the Proxy-Mapper that the data segment object is currently being used by the VM. The Proxy-Mapper must not destroy an open data segment.</p> <p>The <i>KnPxmOpenArgs</i> data structure has the following members:</p> <pre>KnDtPxMapper*  pxm ; KnSgId         sgId ; int            write ;</pre> <p>The <i>KnPxmCloseArgs</i> data structure has exactly the same members. The <i>pxm</i> field is a pointer to the <i>KnExtPxMapper</i> structure previously defined by a call to <i>dcPxmDeclare</i>. It is also associated with the data segment defined by the <i>sgId</i> member of the structure at <i>dcAlloc</i> time. If the <i>write</i> field is set to 0, the data cache will not be modified. If the value is not 0, it may be modified. This allows the Proxy-Mapper to optimize sync or flush operations.</p> <p>The <i>PxmOpen</i> up-call is invoked by the VM in the following cases:</p> <ul style="list-style-type: none"> <li>■ When <i>rgnMapFromDtCache</i> or <i>rgnInitFromDtCache</i> are invoked</li> <li>■ When <i>rgnMapFromActor</i> or <i>rgnInitFromActor</i> or <i>rgnDup</i> are invoked, and the source region was created through a call to <i>rgnMapFromDtCache</i></li> <li>■ When a physical page attached to the data cache will be destroyed by the swapper.</li> </ul> <p>The <i>PxmClose</i> up-call is invoked by the VM in the following cases:</p> <ul style="list-style-type: none"> <li>■ When a region mapping the data segment identified by the <i>sgId</i> field of the <i>KnPxmCloseArgs</i> data structure is destroyed</li> <li>■ When a physical page attached to the corresponding data cache has been destroyed</li> </ul>

**RESTRICTIONS**

The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE**

No error messages are returned.

**ERRORS**

none

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`dcPxmDeclare(2SEG)`, `dcAlloc(2SEG)`, `dcFree(2SEG)`, `svLapCreate(2K)`,  
`lapInvoke(2K)`

<b>NAME</b>	PxmGetAcc – Get access rights and data on a part of a data segment
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; void PxmGetAcc(KnPxmGetAccArgs *getArg, void *cookie);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>PxmGetAcc</i> routine is provided by the Proxy-Mapper and is invoked as an up-call by the CHORUS VM. It is invoked as a <i>lap</i>; its second argument is the <i>lap</i> cookie as defined at <i>svLapCreate</i> time. The <i>PxmGetAcc</i> up-call is invoked by the VM in order to obtain access rights needed to perform an operation on the data cache.</p> <p>The <i>KnPxmGetAccArgs</i> data structure has the following members:</p> <pre>KnDtPxMapper*  pxm ; KnSgId         sgId ; KnPxmAcc*     accParms ; KnPage*       page ; VmSize*       size ;</pre> <p>The <i>pxm</i> field is a pointer to the <i>KnExtPxMapper</i> structure previously defined by a call to <i>dcPxmDeclare</i> and associated with the data segment defined by the <i>sgId</i> member of the structure at <i>dcAlloc</i> time.</p> <p>The access rights are specified by the data structure pointed to by the <i>accParms</i> field. The <i>KnPxmAcc</i> structure contains the following fields:</p> <pre>KnPxmAccReq  req ; KnPxmAccRep  rep ; void*        lock ;</pre> <p>The <i>req</i> field describes the access rights requested by the VM, while the <i>rep</i> field is filled by the Proxy-Mapper to define the actual access rights returned to the VM. The <i>KnPxmAccReq</i> structure contains the following members:</p> <pre>VmFlags  access ; VmOffset start ; VmOffset end ; VmOffset reqStart ; VmOffset reqEnd ;</pre> <p>The <i>access</i> field specifies the access rights, and may be set to either <i>K_READABLE</i> or <i>K_WRITABLE</i>. When set to <i>K_READABLE</i>, the corresponding data will be read—only by the VM. If set to <i>K_WRITABLE</i>, the corresponding data may be modified by the VM. In such a case, the Proxy-Mapper should insure that</p>

backing store (if needed) will be available when data is to be pushed out. The *start* and *end* fields specify the access range desired, the *reqStart* and *reqEnd* specify the access range required.

The *KnPxmAccRep* structure contains the following members:

```
VmFlags    access ;
VmOffset   start ;
VmOffset   end ;
int        diag ;
```

where the *access* field specifies the access rights returned by the Proxy-Mapper to the VM (either *K\_READABLE* or *K\_WRITABLE*). The *start* and *end* fields specify the range of bytes for which the Proxy-Mapper grants access rights to the VM as a result of the *PxmGetAcc* up-call. The *diag* field contains the diagnostic returned by the Proxy-Mapper to the VM. If set to *K\_OK*, the up-call has been successful and the fields returned are meaningful, otherwise it is set to the appropriate error code. If the error code returned is *K\_EBUSY*, the VM will invoke the *PxmGetAcc* up-call again.

The Proxy-Mapper could grant write access in response to a read-only access request, but it must grant write access in response to a write access request. The Proxy-Mapper must guarantee that the [*rep.start*, *rep.end*] range is included in the [*req.start*, *req.end*] range, but also includes the [*req.reqStart*, *req.reqEnd*] range. . The *rep.start* and *rep.end*+1 values must be fragment-aligned.

The VM may request data at the same time as access rights. In this case, the *page* argument specifies the list of pages which should be read from the data segment. The *page* field points to a *KnPage* structure which specifies the first page of the list. The *KnPage* structure has the following members:

```
struct KnPage*  next ;
void*          work ;
```

The Proxy-Mapper must set the *size* output field to the size of valid data (from the beginning of the page list). If no data are requested, the *page* is NULL and the Proxy-Mapper must not refer to the *size* field. The *KnPage* structure is the physical page descriptor exported by the VM. The *next* field points to the next page of the list and is set to NULL for the last page of the list. The *work* field is not used by the VM while the page is owned by the Proxy-Mapper, it may be used by the Proxy-Mapper while it owns the page.

There are two VM operations which can invoke the *PxmGetAcc* up-call:

- *dcRead/dcWrite* invocations
- A page fault on a region created by *rgnMapFromDtCache*.

In the case of *dcRead* or *dcWrite*, the *lock* field contains the *accLock* argument of the *dcRead/dcWrite* call. Thus, the *lock* field is an input argument when *PxmGetAcc* is called as a result of *dcRead* or *dcWrite*.

In the case of a page fault, the *lock* field is set to NULL and the *K\_PAGEFAULT* bit is set in the *req.access* field. If the *PxmGetAcc* is performed successfully, the VM invokes the *PxmRelAccLock* up-call passing the value of the *lock* field as an argument. This value is returned to the VM by the Proxy-Mapper at the completion of the *PxmGetAcc* up-call. Thus, in the case of page fault, the *lock* field is an input/output argument of the *PxmGetAcc* up-call. Note that the *relAccLock* up-call is never invoked as a result of a *dcRead* or *dcWrite* nucleus call.

The pages of the *getArg->page* list are conditionally attached to the data cache. It means that they are not visible to *dcFlush* and therefore *dcFlush* cannot be blocked on them. If *dcFillZero* or *dcTrunc* find a conditionally attached page (which must be filled) they perform the following actions:

- Detach this page from the data cache
- Allocate a new physical page and attach it to the data cache
- Invoke the *PxmPullIn* up-call to read data from the data segment, if needed
- Fill the appropriate part of the page with zero, if needed

Once access rights are obtained, the VM attaches the pages to the data cache fully (they become visible to *dcFlush*) and invokes the *PxmRelAccLock* up-call (if needed, for example in the case of a page fault). Note that, in the case of a pure *PxmGetAcc*, the *getArg->page* is NULL. The *PxmRelAccLock* up-call is invoked when the access rights of the page have already been upgraded.

**RESTRICTIONS**

The current implementation is done only for trusted supervisor actors.

**RETURN VALUE**

None.

**ERRORS**

none

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*dcPxmDeclare(2SEG)*, *dcAlloc(2SEG)*, *svLapCreate(2K)*, *lapInvoke(2K)*, *pageMap(2SEG)*, *dcRead(2SEG)*, *dcWrite(2SEG)*, *rgnMapFromDtCache(2SEG)*, *PxmRelAccLock(2SEG)*

<b>NAME</b>	PxmOpen, PxmClose – Open a Data Segment; Close a Data Segment
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; void PxmOpen(KnPxmOpenArgs * openArg, void * cookie); void PxmClose(KnPxmCloseArgs * closeArgvoid, void * cookie);</pre>
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>PxmOpen/PxmClose</i> routines are provided by the Proxy-Mapper and are invoked as up-calls by the CHORUS VM. They are invoked as <i>laps</i>; their second argument is the lap cookie as defined at <i>svLapCreate</i> time. The <i>PxmOpen/PxmClose</i> up-calls are invoked by the VM in order to inform the Proxy-Mapper that the data segment object is currently being used by the VM. The Proxy-Mapper must not destroy an open data segment.</p> <p>The <i>KnPxmOpenArgs</i> data structure has the following members:</p> <pre>KnDtPxMapper*   pxm ; KnSgId          sgId ; int             write ;</pre> <p>The <i>KnPxmCloseArgs</i> data structure has exactly the same members. The <i>pxm</i> field is a pointer to the <i>KnExtPxMapper</i> structure previously defined by a call to <i>dcPxmDeclare</i>. It is also associated with the data segment defined by the <i>sgId</i> member of the structure at <i>dcAlloc</i> time. If the <i>write</i> field is set to 0, the data cache will not be modified. If the value is not 0, it may be modified. This allows the Proxy-Mapper to optimize sync or flush operations.</p> <p>The <i>PxmOpen</i> up-call is invoked by the VM in the following cases:</p> <ul style="list-style-type: none"> <li>■ When <i>rgnMapFromDtCache</i> or <i>rgnInitFromDtCache</i> are invoked</li> <li>■ When <i>rgnMapFromActor</i> or <i>rgnInitFromActor</i> or <i>rgnDup</i> are invoked, and the source region was created through a call to <i>rgnMapFromDtCache</i></li> <li>■ When a physical page attached to the data cache will be destroyed by the swapper.</li> </ul> <p>The <i>PxmClose</i> up-call is invoked by the VM in the following cases:</p> <ul style="list-style-type: none"> <li>■ When a region mapping the data segment identified by the <i>sgId</i> field of the <i>KnPxmCloseArgs</i> data structure is destroyed</li> <li>■ When a physical page attached to the corresponding data cache has been destroyed</li> </ul>



**RESTRICTIONS**

The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE**

No error messages are returned.

**ERRORS**

none

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`dcPxmDeclare(2SEG)`, `dcAlloc(2SEG)`, `dcFree(2SEG)`, `svLapCreate(2K)`,  
`lapInvoke(2K)`

**NAME**

**SYNOPSIS**

**FEATURES**

**DESCRIPTION**

**RESTRICTIONS**

**RETURN VALUE**

**ATTRIBUTES**

**SEE ALSO**

PxmPullIn – Read data from a data segment

```
#include <mem/chMem.h>
void PxmPullIn(KnPxmPullInArgs *pullArg, void *cookie);
```

PXM\_EXT

**Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.

The *PxmPullIn* routine is provided by the Proxy-Mapper and is invoked as an up-call by the CHORUS VM. It is invoked as a *lap*; its second argument is the lap cookie as defined at *svLapCreate* time. The *PxmPullIn* up-call is invoked by the VM in order to read data from the data segment into the data cache. The *KnPxmPullInArgs* structure has the following members:

```
KnDtPxMapper*   pxm ;
KnSgId          sgId ;
KnPage*        page ;
```

The *sgId* field specifies the corresponding data segment. The *page* field field specifies the list of pages to be read. It points to a *KnPage* structure (see *PxmGetAcc(2SEG)* ) which specifies the first page of the list.

The pages of the *PxmPullIn* page list are fully attached to the data cache and write locked (exclusively). This means that all Nucleus calls (including *dcFlush*, *dcFillZero* and *dcTrunc*) which acquire a lock on this type of page are blocked until *pullIn* is finished.

The VM calls *PxmPullIn* when access rights have already been obtained by a previous *PxmGetAcc*. In the case of page fault, the *PxmRelAccLock* (see the *PxmGetAcc(2SEG)* man page) up-call is invoked after *PxmPullIn*.

The current implementation is only applicable to trusted supervisor actors.

None.

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

*dcPxmDeclare(2SEG)*, *dcAlloc(2SEG)*, *dcFree(2SEG)*, *svLapCreate(2K)*, *lapInvoke(2K)*, *PxmGetAcc(2SEG)*, *PxmRelAccLock(2SEG)*, *pageMap(2SEG)*

<b>NAME</b>	PxmPushOutAsyn – Write asynchronously to a data segment						
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; void PxmPushOutAsyn(KnPxmPushOutArgs *pushArg, void *cookie);</pre>						
<b>FEATURES</b>	PXM_EXT						
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>PxmPushOutAsyn</i> routine is provided by the Proxy-Mapper and is invoked as an up-call by the CHORUS VM. It is invoked as a <i>lap</i>; its second argument is the lap cookie as defined at <i>svLapCreate</i> time. The <i>PxmPushOutAsyn</i> up-call is invoked by the VM in order to write data to a data segment from the data cache. The <i>KnPxmPushOutArgs</i> structure has the following members:</p> <pre>KnDtPxMapper*   pxm ; KnSgId          sgId ; KnPage*        page ; VmFlags        flags ; void*           pdDesc ;</pre> <p>The <i>sgId</i> field specifies the corresponding data segment. The <i>page</i> field specifies the list of pages to be written. It points to a <i>KnPage</i> structure (see <i>PxmGetAcc(2SEG)</i>) which specifies the first page of the list.</p> <p>If at least one page of the list has been modified, The VM sets the K_DIRTY bit in the <i>flags</i> field. The <i>flags</i> field also contains the K_POUT_TYPE bit field which specifies the reason for the <i>PxmPushOutAsyn</i> up-call:</p> <table border="0"> <tr> <td>K_POUT_FLUSH</td> <td>This up-call is the result of a call to <i>dcFlush</i>.</td> </tr> <tr> <td>K_POUT_SYNC</td> <td>This up-call is the result of a call to <i>dcSync</i>.</td> </tr> <tr> <td>K_POUT_SWAP</td> <td>This up-call is invoked by the swapper.</td> </tr> </table> <p>If <i>PxmPushOutAsyn</i> is invoked by <i>dcFlush</i>, the <i>pdDesc</i> field is the <i>pout</i> argument of the <i>dcFlush</i> nucleus call, otherwise it is set to NULL.</p> <p>The Proxy-Mapper must return from the <i>PxmPushOutAsyn</i> up-call and must perform the <i>pageIoDone</i> nucleus call after the write operation is completed. This may be done either before or after returning from the <i>PxmPushOutAsyn</i> up-call. There is no need for the Proxy-Mapper to invoke the <i>pageIoDone</i> call with the same page list. Multiple calls to <i>pageIoDone</i> may be performed one with disjointed sub-lists of pages.</p> <p>The pages of the page list are fully attached to the data cache and locked. In the case of <i>dcFlush</i> or swap, the pages are write locked (exclusively). In the case of <i>dcSync</i>, the pages are read locked (share). Note that in order to avoid a</p>	K_POUT_FLUSH	This up-call is the result of a call to <i>dcFlush</i> .	K_POUT_SYNC	This up-call is the result of a call to <i>dcSync</i> .	K_POUT_SWAP	This up-call is invoked by the swapper.
K_POUT_FLUSH	This up-call is the result of a call to <i>dcFlush</i> .						
K_POUT_SYNC	This up-call is the result of a call to <i>dcSync</i> .						
K_POUT_SWAP	This up-call is invoked by the swapper.						

concurrent *PxmPushOutAsyn* of the same page, *dcSync* first acquires the write lock, and then downgrades it to the read one. This implies that all Nucleus calls (including *dcFlush*, *dcFillZero* and *dcTrunc*) which acquire a lock on this type of page are blocked until *pageIoDone* is finished.

**RESTRICTIONS**

The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE**

No error messages are returned.

**ERRORS**

none

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*dcPxmDeclare(2SEG)*, *dcAlloc(2SEG)*, *dcFree(2SEG)*, *svLapCreate(2K)*, *lapInvoke(2K)*, *PxmGetAcc(2SEG)*, *PxmRelAccLock(2SEG)*, *pageMap(2SEG)*, *pageIoDone(2SEG)*

<b>NAME</b>	PxmRelAccLock – Release an access lock				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; void PxmRelAccLock(KnPxmRelAccArgs *relArg, void *cookie);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>PxmRelAccLock</i> routine is provided by the Proxy-Mapper and is invoked as an up-call by the CHORUS VM. It is invoked as a <i>lap</i>; , its second argument is the lap cookie as defined at <i>svLapCreate</i> time. The <i>PxmRelAccLock</i> up-call is invoked by the VM after a previous up-call to <i>PxmGetAcc</i> resulting from a page fault. In such a case, the Proxy-Mapper must set the output <i>accLock</i> field. When the page fault has been fully processed by the VM, the VM invokes the <i>PxmRelAccLock</i> up-call.</p> <p>The <i>KnPxmRelAccArgs</i> structure has the following members:</p> <pre>KnDtPxMapper*   pxm ; KnSgId          sgId ; void*           lock ;</pre> <p>The <i>sgId</i> identifies the data segment on which the previous <i>PxmGetAcc</i> up-call was performed. The <i>lock</i> field is set to the value returned by the <i>PxmGetAcc</i> up-call.</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	None.				
<b>ERRORS</b>	No error messages are returned.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				
<b>SEE ALSO</b>	<i>dcPxmDeclare(2SEG)</i> , <i>dcAlloc(2SEG)</i> , <i>dcFree(2SEG)</i> , <i>svLapCreate(2K)</i> , <i>lapInvoke(2K)</i> , <i>dcCluster(2SEG)</i>				

**NAME** PxmStat – Get information about a data segment

**SYNOPSIS** #include <mem/chMem.h>  
void PxmStat(KnPxmStatArgs \*statArg, void \*cookie);

**FEATURES** PXM\_EXT

**DESCRIPTION** **Caution** - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.

The *PxmStat* routine is provided by the Proxy-Mapper and is invoked as an up-call by the CHORUS VM. It is invoked as a *lap*; its second argument is the *lap* cookie as defined at *svLapCreate* time. The *PxmStat* up-call is invoked by the VM in order to obtain data segment—specific information as part of a call to *rgnStat*. The *KnPxmStatArgs* structure is composed of the following fields:

```
KnDtPxMapper*   pxm ;
KnSgId          sgId ;
char*           buff ;
int             size ;
```

The *sgId* field specifies the corresponding data segment. The *buff* field points to the buffer to which the information is to be copied. The *size* field specifies the buffer size in bytes. If the buffer size is sufficient for the data segment information to be copied, the Proxy-Mapper copies them to the buffer, otherwise, it does nothing. The *PxmStat* function returns the size of the data segment information.

**RESTRICTIONS** The current implementation is only applicable to trusted supervisor actors.

**RETURN VALUE** If successful, *PxmStat* returns the size of the data segment information in bytes.

**ERRORS** None.

**ATTRIBUTES** See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** *dcPxmDeclare(2SEG)*, *dcAlloc(2SEG)*, *dcFree(2SEG)*, *svLapCreate(2K)*, *lapInvoke(2K)*, *rgnStat(2K)*

<b>NAME</b>	PxmSwapOut – Customize next swap out				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; void PxmSwapOut(KnPxmSwapOutArgs *swArg, void *cookie);</pre>				
<b>FEATURES</b>	PXM_EXT				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>PxmSwapOut</i> routine is provided by the Proxy-Mapper and is invoked as an up-call by the CHORUS VM. It is invoked as a <i>lap</i>; its second argument is the lap cookie as defined at <i>svLapCreate</i> time. The <i>PxmSwapOut</i> up-call is invoked by the VM in order to inform the Proxy-Mapper that a page is going to be pushed out (due to action by the VM daemons), as opposed to a push out resulting from a <i>dcFlush</i> or a <i>dcSync</i>. This up-call is an opportunity for the Proxy-Mapper to tailor the behavior of the VM, so that swap out will occur on ranges that suit the Proxy-Mapper's needs.</p> <p>The <i>KnPxmSwapOutArgs</i> data structure has the following fields:</p> <pre>KnDtPxmMapper*   pxm ; KnSgId           sgId ; KnPage*         page ; VmOffset*       offset ; VmOffset*       end ;</pre> <p>The <i>pxm</i> field points to the Proxy-Mapper definition associated with the segment at the time of <i>dcAlloc</i>. The <i>sgId</i> field specifies the corresponding data segment. The <i>page</i> field points to the list of pages which are going to be pushed out by the VM. The starting offset of this contiguous list of pages is provided at the location pointed to by the <i>offset</i> field. The last offset of the range to be swapped out is defined as the location pointed to by the <i>end</i> field. The <i>PxmSwapOut</i> up-call enables the Proxy-Mapper to change the range of pages to be swapped out, by changing the values pointed to by the input/output fields: <i>offset</i> and <i>end</i>. If no <i>PxmSwapOut</i> up-call routine has been defined in the <i>pxm</i> structure associated with the data segment at <i>dcAlloc</i> time, swap out will be performed by the VM according to its own rules based on the cluster size (see <i>dcCluster</i>).</p>				
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.				
<b>RETURN VALUE</b>	None.				
<b>ERRORS</b>	No error codes are returned.				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				

**SEE ALSO**

dcPxmDeclare(2SEG), dcAlloc(2SEG), dcFree(2SEG), svLapCreate(2K),  
lapInvoke(2K), dcCluster(2SEG)



<b>NAME</b>	rgnFlush – Flush a region
<b>SYNOPSIS</b>	#include <mem/chMem.h> int <b>rgnFlush</b> (KnCap *actorcap, VmOffset start, VmSize size, VmFlags flags);
<b>FEATURES</b>	PXM_EXT
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>rgnFlush</i> call performs a flush operation on the ranges of data caches mapped to an address range of an actor address space. The address range must be page-aligned.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the supervisor address space is used.</p> <p>The <i>flags</i> argument specifies the way the flush is performed as well as the upper access rights for the target parts after the flush operation.</p> <p>If the K_COPYBACK flag is set, <i>rgnFlush</i> writes back all (even clean) cached data of the target parts and does not change the parts' access rights.</p> <p>If the K_WRITABLE flag is set, <i>rgnFlush</i> writes back all modified data of the target parts and does not change the parts' access rights.</p> <p>If the K_READABLE flag is set, <i>rgnFlush</i> writes back all modified data of the target parts and sets the parts' access rights to read only.</p> <p>If the K_FREEZE flag is set, <i>rgnFlush</i> writes back all modified data of the target parts and sets the parts to non-accessible.</p> <p>If the K_NOACCESS flag is set, <i>rgnFlush</i> writes back all modified data of the target parts and invalidates the parts.</p> <p>If the K_DESTROY flag is set, <i>rgnFlush</i> invalidates the target parts without writing back any data.</p> <p>If K_ASYNC flag is set, the vm performs the asynchronous write operations required by <i>rgnFlush</i>.</p> <p>If K_PAGEFAULT flag is set, <i>rgnFlush</i> is the result of a page fault. This type of <i>lcFlush</i> operation could break the atomicity of an <i>sgRead/sgWrite</i> operation running concurrently on the same local cache parts.</p>
<b>RESTRICTIONS</b>	The target actor and the current actor must be located on the same site.
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.
<b>ERRORS</b>	[K_EFAULT]                      Some of the arguments provided are outside the caller's address space.

[K_EINVAL]	An inconsistent actor capability was provided.
[K_EROUND]	<i>start</i> or <i>size</i> isn't page-aligned.
[K_EADDR]	Some or all the addresses from the target address range are invalid.
[K_EOFFSET]	Tried to flush a segment outside the valid offset range in a segment.
[K_EINVAL]	The <i>flags</i> argument contains invalid flag values.
[K_EBUSY]	Tried to invalidate or destroy a no-demand (mapped to a region with the K_NODEMAND attribute) physical memory.

**ATTRIBUTES**

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

`rgnInitFromDtCache(2SEG)`, `rgnAllocate(2K)`, `rgnFree(2K)`

<b>NAME</b>	rgnInit – allocate a region in an actor address space and initialize it from a segment
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnInit(KnCap *actorcap, KnRgnDesc *rgndesc, KnObjDesc *segdesc);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>rgnInit</i> call creates a region in the address space of the target actor and maps a volatile copy of a segment range to the region. The target actor is specified by <i>actorcap</i> - a pointer to the target actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the region is allocated in the supervisor address space and isn't attached to any particular supervisor actor; it cannot therefore be implicitly deallocated by an <i>actorDelete(2K)</i> of a supervisor actor.</p> <p>The <i>rgndesc</i> pointer points to a <i>KnRgnDesc</i> structure containing the specification of the region to be created, as described in <i>rgnAllocate (2K)</i>.</p> <p>The <i>segdesc</i> pointer points to a <i>KnObjDesc</i> structure containing the source data specification. The <i>KnObjDesc</i> fields are as follows:</p> <pre>KnCap      dataObject ; VmOffset   startOffset ; VmSize     size ;</pre> <p>The <i>dataObject</i> field defines the capability of the source segment.</p> <p>The <i>startOffset</i> defines the starting offset of the data in the segment. Its value must be aligned to a virtual page boundary.</p> <p>The <i>size</i> field defines the size of the data.</p> <p>The caller can specify any of the <i>rgndesc-&gt;options</i> as described in <i>rgnAllocate (2K)</i>, except that the K_RESERVED flag is prohibited. The K_FILLZERO option is also interpreted differently here: only the range from <i>rgndesc-&gt;startAddr + segdesc-&gt;size</i> to <i>rgndesc-&gt;startAddr + rgnDesc-&gt;size - 1</i> will be zero-filled.</p> <p>The kernel uses the standard Chorus-IPC based kernel-mapper protocol (see <i>MpPullIn(2SEG)</i>, <i>MpGetAccess(2SEG)</i>) to read the current state of the segment prior to copying it to newly allocated volatile memory, and mapping the memory to the region created. The kernel implementation can defer the effective data reading and copy.</p>
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.

**ERRORS**

- [K\_EFAULT] Some of the arguements provided are outside the caller's address space.
- [K\_EINVAL] An inconsistent actor capability was provided.
- [K\_EUNKNOWN] *actorcap* does not specify a reachable actor.
- [K\_EROUND] *rgndesc->startAddr* is not page-aligned.
- [K\_EROUND] *segdesc->startOffset* is not page-aligned.
- [K\_ESPACE] Tried to create a region outside the valid range for the address space of an actor as returned by *vmStat(2K)*.
- [K\_ESPACE] *rgndesc->size* is zero.
- [K\_EOFFSET] Attempted to map a segment outside its valid offset range as returned by *vmStat(2K)*.
- [K\_ESIZE] *segdesc->size* is greater than *rgndesc->size*.
- [K\_EOVERLAP] The K\_ANYWHERE option was specified and there is insufficient room available in the address space to create the region.
- [K\_EOVERLAP] The K\_RESTRICTIVE option was specified and there is insufficient room in the target address range.
- [K\_EOVERLAP] the region created overlaps an existing region.
- [K\_ENOMEM] The system is out of resources.
- [K\_EMAPPER] The segment mapper doesn't respect the kernel/mapper protocol.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*, *rgnMap(2SEG)*, *rgnStat(2K)*, *MpPullIn(2SEG)*, *vmStat(2K)*

<b>NAME</b>	rgnInitFromDtCache – Create a region in an actor address space and initialize it from a segment										
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnInitFromDtCache(KnCap *actorcap, KnRgnDesc *rgndesc, KnLcId lcId, VmOffset start, VmOffset end);</pre>										
<b>FEATURES</b>	PXM_EXT										
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>rgnInitFromDtCache</i> call allocates a region in the actor address space specified by <i>actorcap</i>. It initializes from the segment associated with the data cache specified by <i>lcId</i>. If <i>actorcap</i> is set to <code>K_MYACTOR</code>, the current actor is used. If <i>actorcap</i> is set to <code>K_SVACTOR</code>, the region is allocated in supervisor address space and is not attached to any particular supervisor actor.</p> <p>The <i>rgnDesc</i> pointer points to a <i>KnRgnDesc</i> structure containing the specification for the region to be created, as described in <i>rgnAllocate(2K)</i>. Note that the <code>K_FILLZERO</code> option is interpreted differently here; the range of bytes comprised between <i>rgndesc-&gt;startAddr+end+1</i> and <i>rgndesc-&gt;startAddr+size-1</i> will be zero-filled.</p> <p>The <i>start</i> argument specifies which offset of the segment will be mapped to the first location of the newly created region. Its value must be aligned on a virtual page boundary. The <i>end</i> argument specifies the offset of the last byte of the data segment to be mapped to the created region.</p>										
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.										
<b>RETURN VALUE</b>	If successful <code>K_OK</code> is returned, otherwise a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td style="padding-right: 20px;">[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td>[K_EINVAL]</td> <td>An inconsistent actor capability was provided.</td> </tr> <tr> <td>[K_ESIZE]</td> <td>The <i>end</i> argument and the <i>size</i> field of the <i>rgndesc</i> argument are inconsistent.</td> </tr> <tr> <td>[K_ENOMEM]</td> <td>The system is out of resources, or some or all of the memory identified by the operation could not be locked when <code>K_NODEMAND</code> and <code>K_NOWAITFORMEMORY</code> have both been specified.</td> </tr> <tr> <td>[K_EOFFSET]</td> <td>Attempted to map a segment outside its valid offset range.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability was provided.	[K_ESIZE]	The <i>end</i> argument and the <i>size</i> field of the <i>rgndesc</i> argument are inconsistent.	[K_ENOMEM]	The system is out of resources, or some or all of the memory identified by the operation could not be locked when <code>K_NODEMAND</code> and <code>K_NOWAITFORMEMORY</code> have both been specified.	[K_EOFFSET]	Attempted to map a segment outside its valid offset range.
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.										
[K_EINVAL]	An inconsistent actor capability was provided.										
[K_ESIZE]	The <i>end</i> argument and the <i>size</i> field of the <i>rgndesc</i> argument are inconsistent.										
[K_ENOMEM]	The system is out of resources, or some or all of the memory identified by the operation could not be locked when <code>K_NODEMAND</code> and <code>K_NOWAITFORMEMORY</code> have both been specified.										
[K_EOFFSET]	Attempted to map a segment outside its valid offset range.										

[K_EOVERLAP]	The K_ANYWHERE option was specified and there is insufficient room in the address space to allocate the region. The K_RETSRICTIVE option was specified and there is insufficient room in the target address range. The region specified overlaps an existing region. The K_RESTRICTIVE option was specified and the right boundary is less than the left one.
[K_EROUND]	The <i>start</i> argument is not a multiple of the page size as returned by <i>vmPageSize(2K)</i> .
[K_ESPACE]	Attempted to allocate a region outside the valid range for the address space of an actor as returned by <i>vmStat(2K)</i>

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnInitFromDtCache(2SEG)*, *rgnAllocate(2K)*, *rgnFree(2K)*

<b>NAME</b>	rgnMap – create a region in an actor address space and map a segment												
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnMap(KnCap *actorcap, KnRgnDesc *rgndesc, KnObjDesc *segdesc);</pre>												
<b>FEATURES</b>	MEM_VM												
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>rgnMap</i> call creates a region in the address space of the target actor and maps a segment to the region.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the target actor capability. If <i>actorcap</i> is K_MYACTOR, the address space of the current actor is used. If <i>actorcap</i> is K_SVACTOR, the region is allocated in the supervisor address space and isn't attached to any particular supervisor actor; it cannot therefore be implicitly deallocated using an <i>actorDelete(2K)</i> of a supervisor actor.</p> <p>The <i>rgndesc</i> pointer points to a <i>KnRgnDesc</i> structure containing the specification of the region to be created as described in <i>rgnAllocate(2K)</i>.</p> <p>The <i>segdesc</i> pointer points to a <i>KnObjDesc</i> structure containing the specification of the segment range to be mapped to the created region described in <i>rgnInit(2SEG)</i>. Note that the <i>size</i> field will be ignored because the size of the range mapped is equal to the size of the region created.</p> <p>The caller can specify any of the <i>rgndesc-&gt;options</i> as described in <i>rgnAllocate(2K)</i> except that the K_FILLZERO flag is ignored, and the K_RESERVED flag is prohibited. Also, the caller can specify K_NOSYNC flag.</p> <p>If the K_NOSYNC flag is specified, the kernel writes back dirty pages of the segment as late as possible. Otherwise, any modification of a segment is written back within a period of time specified using a system configuration parameter..</p>												
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.												
<b>ERRORS</b>	<table border="0"> <tr> <td>[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space.</td> </tr> <tr> <td>[K_EINVAL]</td> <td>An inconsistent actor capability was provided.</td> </tr> <tr> <td>[K_EUNKNOWN]</td> <td><i>actorcap</i> does not specify a reachable actor.</td> </tr> <tr> <td>[K_EROUND]</td> <td><i>rgndesc-&gt;startAddr</i> is not page-aligned.</td> </tr> <tr> <td>[K_EROUND]</td> <td><i>segdesc-&gt;startOffset</i> is not page-aligned.</td> </tr> <tr> <td>[K_ESPACE]</td> <td>Tried to create a region outside the valid range for the address space of an actor as returned by <i>vmStat(2K)</i>.</td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.	[K_EINVAL]	An inconsistent actor capability was provided.	[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.	[K_EROUND]	<i>rgndesc-&gt;startAddr</i> is not page-aligned.	[K_EROUND]	<i>segdesc-&gt;startOffset</i> is not page-aligned.	[K_ESPACE]	Tried to create a region outside the valid range for the address space of an actor as returned by <i>vmStat(2K)</i> .
[K_EFAULT]	Some of the arguments provided are outside the caller's address space.												
[K_EINVAL]	An inconsistent actor capability was provided.												
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.												
[K_EROUND]	<i>rgndesc-&gt;startAddr</i> is not page-aligned.												
[K_EROUND]	<i>segdesc-&gt;startOffset</i> is not page-aligned.												
[K_ESPACE]	Tried to create a region outside the valid range for the address space of an actor as returned by <i>vmStat(2K)</i> .												

[K_ESPACE]	<i>rgndesc-&gt;size</i> is zero.
[K_EOFFSET]	Attempted to map a segment outside its valid offset range as returned by <i>vmStat(2K)</i> .
[K_EOVERLAP]	The K_ANYWHERE option was specified and there is insufficient room available in the address space to create the region.
[K_EOVERLAP]	The K_RESTRICTIVE option was specified and there is insufficient room in the target address range.
[K_EOVERLAP]	The region created overlaps an existing region.
[K_ENOMEM]	The system is out of resources.
[K_EMAPPER]	The segment mapper doesn't respect the kernel/mapper protocol.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*rgnAllocate(2K)*, *rgnInit(2SEG)*, *rgnStat(2K)*, *MpPullIn(2SEG)*, *MpGetAccess(2SEG)*



<b>NAME</b>	rgnMapFromDtCache – Create a region in an actor address space and map a segment										
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int rgnMapFromDtCache(KnCap *actorcap, KnRgnDesc *rgndesc, KnLcId lcId, VmOffset start);</pre>										
<b>FEATURES</b>	PXM_EXT										
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>rgnMapFromDtCache</i> call creates a region in the actor address space specified by <i>actorcap</i>, and maps the segment associated with the data cache specified by <i>lcId</i> to it. If <i>actorcap</i> is set to <code>K_MYACTOR</code>, the current actor is used. If <i>actorcap</i> is set to <code>K_SVACTOR</code>, the region is allocated in the supervisor address space and is not attached to any particular supervisor actor.</p> <p>The <i>rgnDesc</i> pointer points to a <i>KnRgnDesc</i> structure containing the specification for the region to be created as described in <i>rgnAllocate(2K)</i>. Note that the <code>K_FILLZERO</code> option will be ignored and that the <code>K_RESERVED</code> flag is prohibited.</p> <p>The <i>start</i> argument specifies which offset of the segment will be mapped to the first location of the newly created region. Its value must be aligned on a virtual page boundary.</p>										
<b>RESTRICTIONS</b>	The current implementation is only applicable to trusted supervisor actors.										
<b>RETURN VALUE</b>	If successful <code>K_OK</code> is returned, otherwise a negative error code is returned.										
<b>ERRORS</b>	<table border="0"> <tr> <td style="vertical-align: top;">[K_EFAULT]</td> <td>Some of the arguments provided are outside the caller's address space. This error will only be returned if the invoker is a user actor.</td> </tr> <tr> <td style="vertical-align: top;">[K_EINVAL]</td> <td>An inconsistent actor capability was given.</td> </tr> <tr> <td style="vertical-align: top;">[K_ENOMEM]</td> <td>The system is out of resources or some or all of the memory identified by the operation could not be locked when <code>K_NODEMAND</code> and <code>K_NOWAITFORMEMORY</code> were both specified.</td> </tr> <tr> <td style="vertical-align: top;">[K_EOFFSET]</td> <td>Attempted to map a segment outside its valid offset range.</td> </tr> <tr> <td style="vertical-align: top;">[K_EOVERLAP]</td> <td>The <code>K_ANYWHERE</code> option was specified and there is insufficient room in the address space to allocate the region. The <code>K_RESTRICTIVE</code> option was specified and there is insufficient room in the target address range. The specified region overlaps an existing region. The <code>K_RESTRICTIVE</code></td> </tr> </table>	[K_EFAULT]	Some of the arguments provided are outside the caller's address space. This error will only be returned if the invoker is a user actor.	[K_EINVAL]	An inconsistent actor capability was given.	[K_ENOMEM]	The system is out of resources or some or all of the memory identified by the operation could not be locked when <code>K_NODEMAND</code> and <code>K_NOWAITFORMEMORY</code> were both specified.	[K_EOFFSET]	Attempted to map a segment outside its valid offset range.	[K_EOVERLAP]	The <code>K_ANYWHERE</code> option was specified and there is insufficient room in the address space to allocate the region. The <code>K_RESTRICTIVE</code> option was specified and there is insufficient room in the target address range. The specified region overlaps an existing region. The <code>K_RESTRICTIVE</code>
[K_EFAULT]	Some of the arguments provided are outside the caller's address space. This error will only be returned if the invoker is a user actor.										
[K_EINVAL]	An inconsistent actor capability was given.										
[K_ENOMEM]	The system is out of resources or some or all of the memory identified by the operation could not be locked when <code>K_NODEMAND</code> and <code>K_NOWAITFORMEMORY</code> were both specified.										
[K_EOFFSET]	Attempted to map a segment outside its valid offset range.										
[K_EOVERLAP]	The <code>K_ANYWHERE</code> option was specified and there is insufficient room in the address space to allocate the region. The <code>K_RESTRICTIVE</code> option was specified and there is insufficient room in the target address range. The specified region overlaps an existing region. The <code>K_RESTRICTIVE</code>										

option was specified and the right boundary is less than the left one.

[K\_EROUND]

The *start* argument is not a multiple of the page size as returned by *vmPageSize(2K)*.

[K\_ESPACE]

Attempted to allocate a region outside the valid range for the address space of an actor as returned by *vmStat(2K)*

#### ATTRIBUTES

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

#### SEE ALSO

*rgnInitFromDtCache(2SEG)*, *rgnAllocate(2K)*, *rgnFree(2K)*

<b>NAME</b>	lcFlush, lcSetRights, sgFlush, sgSyncAll, vmFlush – flush local cache(s)
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcFlush(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int lcSetRights(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int sgFlush(KnObjDesc * segdesc, VmFlags flags);  int sgSyncAll(void);  int vmFlush(KnCap * actorcap, VmAddr address, VmSize size, VmFlags flags);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcFlush</i> call performs a flush operation on a range of a local cache. The <i>lcdesc</i> argument points to a <i>KnObjDesc</i> structure whose members are the following:</p> <pre>KnCap      dataObject ; VmOffset   startOffset ; VmSize     size ;</pre> <p>where the <i>dataObject</i> field is the local cache capability, the <i>startOffset</i> field is the range start offset in the local cache, and the <i>size</i> field is the range size. Both <i>startOffset</i> and <i>size</i> must be fragment-aligned. The size of the fragment is implementation-dependent and is usually equal to the virtual page size divided by 8 (number of bits in one byte).</p> <p>For <i>lcFlush</i>, the target local cache is directly specified by its capability; the flush operation is applied even if the cache is managed by a remote site.</p> <p>The <i>sgFlush</i> call takes the same arguments, except the <i>dataObject</i> field specifies the capability of the cached segment. For <i>sgFlush</i>, the target local cache is indirectly specified by the capability of the corresponding segment; the flush operation implicitly applies to the segment's cache(s) located on the caller's site.</p> <p>The <i>vmFlush</i> call performs a flush operation on the ranges of the local caches mapped to an address range of an actor address space. The address range must be page-aligned.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <i>K_MYACTOR</i>, the address space of the current actor is used. If <i>actorcap</i> is <i>K_SVACTOR</i>, the supervisor address space is used.</p> <p>The <i>flags</i> argument specifies the mode of the flush and the upper access rights for the target parts after the flush operation.</p>

If the `K_COPYBACK` flag is set, *xxFlush* writes back any (even clean) cached data of the target parts and keeps the parts access rights unchanged.

If the `K_WRITABLE` flag is set, *xxFlush* writes back all modified data of the target parts and keeps the parts access rights unchanged.

If the `K_READABLE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts to read access only.

If the `K_FREEZE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts as non-accessible.

If the `K_NOACCESS` flag is set, *xxFlush* writes back all modified data of the target parts and invalidates them.

If the `K_DESTROY` flag is set, *xxFlush* invalidates the target parts without writing back.

If the `K_ASYNC` flag is set, the vm performs the write operations required by the *xxFlush* asynchronously.

If `K_PAGEFAULT` flag is set, *xxFlush* is a result of a page fault. This type of *lcFlush* operation could break the atomicity of a *sgRead*/*sgWrite* operation running concurrently on the same local cache parts.

The `K_ORDERED` flag can only be used with an *lcFlush* or an *lcSetRights* request. If the flag is set, the *orderNb* argument specifies the *lcFlush* message order number. It allows the kernel to detect the situation when two messages sent by a mapper in one order are received by the kernel in another. For instance, if the mapper grants certain access rights in a *MpGetAccess* reply first, then, processing another request, calls an *lcFlush* to recall the access rights, considering the order numbers, the kernel is aware that the access returned by the *MpGetAccess* reply was already recalled by the mapper, and performs another *MpGetAccess* request.

The *lcSetRights* operation is similar to the *lcFlush* operation, except that it only changes data protections without invalidation and/or writing back: the `K_WRITABLE` flag is ignored, the `K_READABLE` flag sets the target parts read—only, whereas the `K_FREEZE` and `K_NOACCESS` flags set the target parts to non-accessible. The `K_COPYBACK`, `K_ASYNC` and `K_DESTROY` flags are prohibited.

The *sgSyncAll* operation writes back asynchronously all dirty parts of all local caches on the site, except the local caches mapped at least once to a region with the `K_NOSYNC` attribute (see *rgnMap(2SEG)*).

## RETURN VALUE

If successful `K_OK` is returned, otherwise a negative error code is returned.

## ERRORS

[`K_EFAULT`] Some of the arguments provided are outside the caller's address space.

[K_EINVAL]	An inconsistent actor capability was provided.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EROUND]	<i>address</i> or <i>size</i> isn't page-aligned.
[K_EROUND]	<i>lcdesc-&gt;startOffset</i> or <i>lcdesc-&gt;size</i> isn't fragment-aligned.
[K_EROUND]	<i>segdesc-&gt;startOffset</i> or <i>segdesc-&gt;size</i> isn't fragment-aligned.
[K_EADDR]	Some or all the addresses from the target address range are invalid.
[K_EOFFSET]	Tried to flush a segment outside the valid offset range in a segment, as returned by <i>vmStat</i> .
[K_EINVAL]	The <i>flags</i> argument contains invalid flag values.
[K_EBUSY]	Tried to invalidate or destroy a no-demand (mapped to a region with K_NODEMAND attribute) physical memory.
[K_EFAIL]	An <i>ipcCall</i> transaction failed during the remote <i>lcFlush</i> or <i>lcSetRights</i> .
[K_EMAPPER]	The mapper doesn't respect the vm/mapper protocol.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site ( *vmFlush* only).

The *sgFlush* and *sgSyncAll* calls remain in the interface for backward compatibility. They will be removed in a future release.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpGetAccess(2SEG)* , *MpPullIn(2SEG)* , *MpPushOut(2SEG)* , *rgnMap(2SEG)* , *lcOpen(2SEG)* , *lcFillZero(2SEG)* , *lcTrunc(2SEG)*

<b>NAME</b>	sgRead, sgWrite – Read data from a segment into an actor address space; Write data from an actor address space into a segment
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int sgRead(KnObjDesc * srcsegdesc, KnCap * dstactorcap, VmAddr dstaddress, VmFlags flags);  int sgWrite(KnObjDesc * dstsegdesc, KnCap * srcactorcap, VmAddr srcaddress, VmFlags flags);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>sgRead</i> function reads data from a segment into an actor address space. The source data is specified by the <i>srcsegdesc</i> argument, which points to a <i>KnObjDesc</i> structure described in <i>rgnInit(2SEG)</i> . The destination is specified by <i>dstactorcap</i> - the destination actor capability, and <i>dstaddress</i> - the start address in the destination actor address space.</p> <p>The <i>sgWrite</i> function writes data from an actor address space into a segment. The source data is specified by <i>srcactorcap</i> - the source actor capability, and <i>srcaddress</i> - the data start address in the source actor address space. The destination is specified by the <i>dstsegdesc</i> argument which points to a <i>KnObjDesc</i> structure described in <i>rgnInit(2SEG)</i> .</p> <p>If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_MYACTOR, the current actor is used. If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_SUPERVISOR, the supervisor address space is used.</p> <p>The <i>flags</i> argument is a combination of the following options:</p> <p><b>K_MOVE</b> This option indicates that after the operation the contents of the source, from the start to the end of the target data, may be undefined. In this case, the system will try to remap the physical pages containing the data rather than copying them.</p> <p><b>K_MOVEAL</b> This option indicates that even if the target data size is not aligned on a page boundary (see <i>vmPageSize(2K)</i> ), the last page partially covered by the data may be remapped. This means that after the operation the contents of the source, from the start to the next page boundary following the end of the target data, and the contents of the destination, from the end of the target data to the next page boundary, may be undefined.</p> <p>If any error occurs during an <i>sgRead</i> or <i>sgWrite</i> operation, the number of bytes read or written before the error is returned in the <code>size</code> field of the corresponding <i>KnObjDesc</i> structure.</p>

**RETURN VALUE** If successful K\_OK is returned, otherwise a negative error code is returned.

<b>ERRORS</b>	[K_EADDR]	The address is out of any allocated region.
	[K_EFAULT]	Some of the arguments provided are outside the caller's address space.
	[K_EINVAL]	An inconsistent actor capability was given.
	[K_ESIZE]	The <i>startOffset</i> and <i>size</i> fields of the <i>KnObjDesc</i> structure are inconsistent.
	[K_EOFFSET]	Attempted to map a segment outside its valid offset range returned by <i>vmStat</i> .
	[K_EPROT]	Attempted to write to a read only region.
	[K_EUNKNOWN]	<i>srcactorcap</i> or <i>dstactorcap</i> does not specify a reachable actor.

**RESTRICTIONS** The target actor, the source actor and the current actor must be located on the same site.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `rgnInit(2SEG)`

<b>NAME</b>	lcStat, sgStat – get the statistics of a local cache				
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; #include &lt;mem/chMapper.h&gt; int lcStat(KnCap * lccap, KnLcStat * stat);  int sgStat(KnCap * segcap, KnLcStat * stat);</pre>				
<b>FEATURES</b>	MEM_VM				
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcStat</i> and <i>sgStat</i> functions get the statistics of the local cache specified either by <i>lccap</i> - a local cache capability, or by <i>segcap</i> - a segment capability. In the latter case, the statistics of the cache of the segment on the current site are returned.</p> <p>The <i>KnLcStat</i> structure describes the statistics associated with a local cache, as follows:</p> <pre>VmSize  physMem ; VmSize  lockMem ; KnCap   segcap  ; KnCap   lccap   ;</pre> <p>The <i>physMem</i> field specifies the physical memory size currently allocated for the local cache.</p> <p>The <i>lockMem</i> field specifies the physical memory size currently fixed for the local cache.</p> <p>The <i>segcap</i> field specifies the capability of the corresponding segment. It is returned by the <i>lcStat</i> call only when the caller is a system actor or when the current thread executes in privileged mode.</p> <p>The <i>lccap</i> field specifies the capability of the segment's local cache on the current site. It is returned by the <i>sgStat</i> call only when the caller is a system actor or when the current thread executes in privileged mode.</p>				
<b>RETURN VALUE</b>	If successful K_OK is returned, otherwise a negative error code is returned.				
<b>ERRORS</b>	<p>[K_EFAULT]                   Some of the arguments provided are outside the caller's address space.</p> <p>[K_EUNDEF]                   The segment specified is not cached on the site.</p>				
<b>ATTRIBUTES</b>	See <i>attributes(5)</i> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Interface Stability</td> <td style="text-align: center;">Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Evolving				



**SEE ALSO**

lcFlush(2SEG)

<b>NAME</b>	lcFlush, lcSetRights, sgFlush, sgSyncAll, vmFlush – flush local cache(s)
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcFlush(KnObjDesc * ldesc, VmFlags flags, unsigned long ordernb);  int lcSetRights(KnObjDesc * ldesc, VmFlags flags, unsigned long ordernb);  int sgFlush(KnObjDesc * segdesc, VmFlags flags);  int sgSyncAll(void);  int vmFlush(KnCap * actorcap, VmAddr address, VmSize size, VmFlags flags);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcFlush</i> call performs a flush operation on a range of a local cache. The <i>ldesc</i> argument points to a <i>KnObjDesc</i> structure whose members are the following:</p> <pre>KnCap      dataObject ; VmOffset   startOffset ; VmSize     size ;</pre> <p>where the <i>dataObject</i> field is the local cache capability, the <i>startOffset</i> field is the range start offset in the local cache, and the <i>size</i> field is the range size. Both <i>startOffset</i> and <i>size</i> must be fragment-aligned. The size of the fragment is implementation-dependent and is usually equal to the virtual page size divided by 8 (number of bits in one byte).</p> <p>For <i>lcFlush</i>, the target local cache is directly specified by its capability; the flush operation is applied even if the cache is managed by a remote site.</p> <p>The <i>sgFlush</i> call takes the same arguments, except the <i>dataObject</i> field specifies the capability of the cached segment. For <i>sgFlush</i>, the target local cache is indirectly specified by the capability of the corresponding segment; the flush operation implicitly applies to the segment's cache(s) located on the caller's site.</p> <p>The <i>vmFlush</i> call performs a flush operation on the ranges of the local caches mapped to an address range of an actor address space. The address range must be page-aligned.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <i>K_MYACTOR</i>, the address space of the current actor is used. If <i>actorcap</i> is <i>K_SVACTOR</i>, the supervisor address space is used.</p> <p>The <i>flags</i> argument specifies the mode of the flush and the upper access rights for the target parts after the flush operation.</p>

If the `K_COPYBACK` flag is set, *xxFlush* writes back any (even clean) cached data of the target parts and keeps the parts access rights unchanged.

If the `K_WRITABLE` flag is set, *xxFlush* writes back all modified data of the target parts and keeps the parts access rights unchanged.

If the `K_READABLE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts to read access only.

If the `K_FREEZE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts as non-accessible.

If the `K_NOACCESS` flag is set, *xxFlush* writes back all modified data of the target parts and invalidates them.

If the `K_DESTROY` flag is set, *xxFlush* invalidates the target parts without writing back.

If the `K_ASYNC` flag is set, the vm performs the write operations required by the *xxFlush* asynchronously.

If `K_PAGEFAULT` flag is set, *xxFlush* is a result of a page fault. This type of *lcFlush* operation could break the atomicity of a *sgRead*/*sgWrite* operation running concurrently on the same local cache parts.

The `K_ORDERED` flag can only be used with an *lcFlush* or an *lcSetRights* request. If the flag is set, the *orderNb* argument specifies the *lcFlush* message order number. It allows the kernel to detect the situation when two messages sent by a mapper in one order are received by the kernel in another. For instance, if the mapper grants certain access rights in a *MpGetAccess* reply first, then, processing another request, calls an *lcFlush* to recall the access rights, considering the order numbers, the kernel is aware that the access returned by the *MpGetAccess* reply was already recalled by the mapper, and performs another *MpGetAccess* request.

The *lcSetRights* operation is similar to the *lcFlush* operation, except that it only changes data protections without invalidation and/or writing back: the `K_WRITABLE` flag is ignored, the `K_READABLE` flag sets the target parts read—only, whereas the `K_FREEZE` and `K_NOACCESS` flags set the target parts to non-accessible. The `K_COPYBACK`, `K_ASYNC` and `K_DESTROY` flags are prohibited.

The *sgSyncAll* operation writes back asynchronously all dirty parts of all local caches on the site, except the local caches mapped at least once to a region with the `K_NOSYNC` attribute (see *rgnMap(2SEG)*).

## RETURN VALUE

If successful `K_OK` is returned, otherwise a negative error code is returned.

## ERRORS

[`K_EFAULT`] Some of the arguments provided are outside the caller's address space.

- [K\_EINVAL] An inconsistent actor capability was provided.
- [K\_EUNKNOWN] *actorcap* does not specify a reachable actor.
- [K\_EROUND] *address* or *size* isn't page-aligned.
- [K\_EROUND] *lcdesc->startOffset* or *lcdesc->size* isn't fragment-aligned.
- [K\_EROUND] *segdesc->startOffset* or *segdesc->size* isn't fragment-aligned.
- [K\_EADDR] Some or all the addresses from the target address range are invalid.
- [K\_EOFFSET] Tried to flush a segment outside the valid offset range in a segment, as returned by *vmStat*.
- [K\_EINVAL] The *flags* argument contains invalid flag values.
- [K\_EBUSY] Tried to invalidate or destroy a no-demand (mapped to a region with K\_NODEMAND attribute) physical memory.
- [K\_EFAIL] An *ipcCall* transaction failed during the remote *lcFlush* or *lcSetRights*.
- [K\_EMAPPER] The mapper doesn't respect the vm/mapper protocol.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site ( *vmFlush* only).  
 The *sgFlush* and *sgSyncAll* calls remain in the interface for backward compatibility. They will be removed in a future release.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpGetAccess(2SEG)* , *MpPullIn(2SEG)* , *MpPushOut(2SEG)* , *rgnMap(2SEG)* , *lcOpen(2SEG)* , *lcFillZero(2SEG)* , *lcTrunc(2SEG)*

<b>NAME</b>	sgRead, sgWrite – Read data from a segment into an actor address space; Write data from an actor address space into a segment
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int sgRead(KnObjDesc * srcsegdesc, KnCap * dstactorcap, VmAddr dstaddress, VmFlags flags);  int sgWrite(KnObjDesc * dstsegdesc, KnCap * srcactorcap, VmAddr srcaddress, VmFlags flags);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>sgRead</i> function reads data from a segment into an actor address space. The source data is specified by the <i>srcsegdesc</i> argument, which points to a <i>KnObjDesc</i> structure described in <i>rgnInit(2SEG)</i> . The destination is specified by <i>dstactorcap</i> - the destination actor capability, and <i>dstaddress</i> - the start address in the destination actor address space.</p> <p>The <i>sgWrite</i> function writes data from an actor address space into a segment. The source data is specified by <i>srcactorcap</i> - the source actor capability, and <i>srcaddress</i> - the data start address in the source actor address space. The destination is specified by the <i>dstsegdesc</i> argument which points to a <i>KnObjDesc</i> structure described in <i>rgnInit(2SEG)</i> .</p> <p>If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_MYACTOR, the current actor is used. If <i>srcactorcap</i> and/or <i>dstactorcap</i> is K_SUPERVISOR, the supervisor address space is used.</p> <p>The <i>flags</i> argument is a combination of the following options:</p> <p><b>K_MOVE</b> This option indicates that after the operation the contents of the source, from the start to the end of the target data, may be undefined. In this case, the system will try to remap the physical pages containing the data rather than copying them.</p> <p><b>K_MOVEAL</b> This option indicates that even if the target data size is not aligned on a page boundary (see <i>vmPageSize(2K)</i> ), the last page partially covered by the data may be remapped. This means that after the operation the contents of the source, from the start to the next page boundary following the end of the target data, and the contents of the destination, from the end of the target data to the next page boundary, may be undefined.</p> <p>If any error occurs during an <i>sgRead</i> or <i>sgWrite</i> operation, the number of bytes read or written before the error is returned in the <i>size</i> field of the corresponding <i>KnObjDesc</i> structure.</p>

**RETURN VALUE** If successful K\_OK is returned, otherwise a negative error code is returned.

- ERRORS**
- [K\_EADDR] The address is out of any allocated region.
  - [K\_EFAULT] Some of the arguments provided are outside the caller's address space.
  - [K\_EINVAL] An inconsistent actor capability was given.
  - [K\_ESIZE] The *startOffset* and *size* fields of the *KnObjDesc* structure are inconsistent.
  - [K\_EOFFSET] Attempted to map a segment outside its valid offset range returned by *vmStat*.
  - [K\_EPROT] Attempted to write to a read only region.
  - [K\_EUNKNOWN] *srcactorcap* or *dstactorcap* does not specify a reachable actor.

**RESTRICTIONS** The target actor, the source actor and the current actor must be located on the same site.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO** `rgnInit(2SEG)`

<b>NAME</b>	lcFlush, lcSetRights, sgFlush, sgSyncAll, vmFlush – flush local cache(s)
<b>SYNOPSIS</b>	<pre>#include &lt;mem/chMem.h&gt; int lcFlush(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int lcSetRights(KnObjDesc * lcdesc, VmFlags flags, unsigned long ordernb);  int sgFlush(KnObjDesc * segdesc, VmFlags flags);  int sgSyncAll(void);  int vmFlush(KnCap * actorcap, VmAddr address, VmSize size, VmFlags flags);</pre>
<b>FEATURES</b>	MEM_VM
<b>DESCRIPTION</b>	<p><b>Caution</b> - This system call is strictly reserved for internal use only. It MUST NOT be used by any application.</p> <p>The <i>lcFlush</i> call performs a flush operation on a range of a local cache. The <i>lcdesc</i> argument points to a <i>KnObjDesc</i> structure whose members are the following:</p> <pre>KnCap      dataObject ; VmOffset   startOffset ; VmSize     size ;</pre> <p>where the <i>dataObject</i> field is the local cache capability, the <i>startOffset</i> field is the range start offset in the local cache, and the <i>size</i> field is the range size. Both <i>startOffset</i> and <i>size</i> must be fragment-aligned. The size of the fragment is implementation-dependent and is usually equal to the virtual page size divided by 8 (number of bits in one byte).</p> <p>For <i>lcFlush</i>, the target local cache is directly specified by its capability; the flush operation is applied even if the cache is managed by a remote site.</p> <p>The <i>sgFlush</i> call takes the same arguments, except the <i>dataObject</i> field specifies the capability of the cached segment. For <i>sgFlush</i>, the target local cache is indirectly specified by the capability of the corresponding segment; the flush operation implicitly applies to the segment's cache(s) located on the caller's site.</p> <p>The <i>vmFlush</i> call performs a flush operation on the ranges of the local caches mapped to an address range of an actor address space. The address range must be page-aligned.</p> <p>The target actor is specified by <i>actorcap</i> - a pointer to the actor capability. If <i>actorcap</i> is <i>K_MYACTOR</i>, the address space of the current actor is used. If <i>actorcap</i> is <i>K_SVACTOR</i>, the supervisor address space is used.</p> <p>The <i>flags</i> argument specifies the mode of the flush and the upper access rights for the target parts after the flush operation.</p>

If the `K_COPYBACK` flag is set, *xxFlush* writes back any (even clean) cached data of the target parts and keeps the parts access rights unchanged.

If the `K_WRITABLE` flag is set, *xxFlush* writes back all modified data of the target parts and keeps the parts access rights unchanged.

If the `K_READABLE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts to read access only.

If the `K_FREEZE` flag is set, *xxFlush* writes back all modified data of the target parts and then sets the parts as non-accessible.

If the `K_NOACCESS` flag is set, *xxFlush* writes back all modified data of the target parts and invalidates them.

If the `K_DESTROY` flag is set, *xxFlush* invalidates the target parts without writing back.

If the `K_ASYNC` flag is set, the vm performs the write operations required by the *xxFlush* asynchronously.

If `K_PAGEFAULT` flag is set, *xxFlush* is a result of a page fault. This type of *lcFlush* operation could break the atomicity of a *sgRead*/*sgWrite* operation running concurrently on the same local cache parts.

The `K_ORDERED` flag can only be used with an *lcFlush* or an *lcSetRights* request. If the flag is set, the *orderNb* argument specifies the *lcFlush* message order number. It allows the kernel to detect the situation when two messages sent by a mapper in one order are received by the kernel in another. For instance, if the mapper grants certain access rights in a *MpGetAccess* reply first, then, processing another request, calls an *lcFlush* to recall the access rights, considering the order numbers, the kernel is aware that the access returned by the *MpGetAccess* reply was already recalled by the mapper, and performs another *MpGetAccess* request.

The *lcSetRights* operation is similar to the *lcFlush* operation, except that it only changes data protections without invalidation and/or writing back: the `K_WRITABLE` flag is ignored, the `K_READABLE` flag sets the target parts read—only, whereas the `K_FREEZE` and `K_NOACCESS` flags set the target parts to non-accessible. The `K_COPYBACK`, `K_ASYNC` and `K_DESTROY` flags are prohibited.

The *sgSyncAll* operation writes back asynchronously all dirty parts of all local caches on the site, except the local caches mapped at least once to a region with the `K_NOSYNC` attribute (see *rgnMap(2SEG)*).

## RETURN VALUE

If successful `K_OK` is returned, otherwise a negative error code is returned.

## ERRORS

[`K_EFAULT`] Some of the arguments provided are outside the caller's address space.



[K_EINVAL]	An inconsistent actor capability was provided.
[K_EUNKNOWN]	<i>actorcap</i> does not specify a reachable actor.
[K_EROUND]	<i>address</i> or <i>size</i> isn't page-aligned.
[K_EROUND]	<i>lcdesc-&gt;startOffset</i> or <i>lcdesc-&gt;size</i> isn't fragment-aligned.
[K_EROUND]	<i>segdesc-&gt;startOffset</i> or <i>segdesc-&gt;size</i> isn't fragment-aligned.
[K_EADDR]	Some or all the addresses from the target address range are invalid.
[K_EOFFSET]	Tried to flush a segment outside the valid offset range in a segment, as returned by <i>vmStat</i> .
[K_EINVAL]	The <i>flags</i> argument contains invalid flag values.
[K_EBUSY]	Tried to invalidate or destroy a no-demand (mapped to a region with K_NODEMAND attribute) physical memory.
[K_EFAIL]	An <i>ipcCall</i> transaction failed during the remote <i>lcFlush</i> or <i>lcSetRights</i> .
[K_EMAPPER]	The mapper doesn't respect the vm/mapper protocol.

**RESTRICTIONS**

The target actor and the current actor must be located on the same site ( *vmFlush* only).

The *sgFlush* and *sgSyncAll* calls remain in the interface for backward compatibility. They will be removed in a future release.

**ATTRIBUTES**

See *attributes(5)* for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**SEE ALSO**

*MpGetAccess(2SEG)* , *MpPullIn(2SEG)* , *MpPushOut(2SEG)* , *rgnMap(2SEG)* , *lcOpen(2SEG)* , *lcFillZero(2SEG)* , *lcTrunc(2SEG)*



# Index

---

## D

**dcAlloc** — Allocate a data cache for a segment; Free a previously allocated data cache 13, 20

**dcCluster** — Set the input and output cluster sizes of a data cache 15

**dcFillZero** — Fill a data segment with zero 16

**dcFlush** — Sync a data cache object; Flush a data cache object 17, 29

**dcFree** — Allocate a data cache for a segment; Free a previously allocated data cache 13, 20

**dcGetPages** — Get a list of pages for read-ahead purpose 22

**dcIsDirty** — Test and reset data cache dirty bit 24

**dcPgNumber** — Get the number of pages for a data cache object 25

**dcPxmDeclare** — Initialize an external proxy-mapper descriptor 26

**dcRead** — Read data from a data cache; Write data to a data cache 27, 33

**dcSync** — Sync a data cache object; Flush a data cache object 17, 29

**dcTrunc** — Truncate a data segment 32

**dcWrite** — Read data from a data cache; Write data to a data cache 27, 33

## L

**lcCap** — Find or create a local cache object for a segment; Release a local cache

object; Return the cabability of a local cache 35–36, 42

**lcClose** — Find or create a local cache object for a segment; Release a local cache object; Return the cabability of a local cache 35–36, 42

**lcFillZero** — zero a range of a local cache 37

**lcFlush** — flush local cache(s) 39, 47, 99, 106, 111

**lcOpen** — Find or create a local cache object for a segment; Release a local cache object; Return the cabability of a local cache 35–36, 42

**lcPushData** — push data from a source local cache to a target local cache 43

**lcRead** — Read data through a local cache into an actor address space; Write data from an actor address space through a local cache 45, 54

**lcSetRights** — flush local cache(s) 39, 47, 99, 106, 111

**lcStat** — get the statistics of a local cache 50, 104

**lcTrunc** — shape the end of a local cache 52

**lcWrite** — Read data through a local cache into an actor address space; Write data from an actor address space through a local cache 45, 54

## M

**MpCreate** — create a temporary segment at the default mapper 56  
**MpGetAccess** — get access to data through a local cache 57  
**MpPullIn** — read data into a local cache 61  
**MpPushOut** — write data back to mapper 64  
**MpRelease** — release a temporary segment or notify a local cache destruction 66

## P

**pageIoDone** — Inform nucleus I/O is complete on a page list 67  
**pageMap** — Map a list of pages in the current actor address space; Unmap a list of pages 68, 73  
**pagePhysAddr** — Get the physical address of a page 70  
**pageSetDirty** — Mark a page as dirty 71  
**pageSgId** — Get the segment identifier associated with a page 72  
**pageUnmap** — Map a list of pages in the current actor address space; Unmap a list of pages 68, 73  
**PxmClose** — Open a Data Segment; Close a Data Segment 75, 80  
**PxmGetAcc** — Get access rights and data on a part of a data segment 77  
**PxmOpen** — Open a Data Segment; Close a Data Segment 75, 80  
**PxmPullIn** — Read data from a data segment 82  
**PxmPushOutAsyn** — Write asynchronously to a data segment 83  
**PxmRelAccLock** — Release an access lock 85  
**PxmStat** — Get information about a data segment 86

**PxmSwapOut** — Customize next swap out 87

## R

**rgnFlush** — Flush a region 89  
**rgnInit** — allocate a region in an actor address space and initialize it from a segment 91  
**rgnInitFromDtCache** — Create a region in an actor address space and initialize it from a segment 93  
**rgnMap** — create a region in an actor address space and map a segment 95  
**rgnMapFromDtCache** — Create a region in an actor address space and map a segment 97

## S

**sgFlush** — flush local cache(s) 39, 47, 99, 106, 111  
**sgRead** — Read data from a segment into an actor address space; Write data from an actor address space into a segment 102, 109  
**sgStat** — get the statistics of a local cache 50, 104  
**sgSyncAll** — flush local cache(s) 39, 47, 99, 106, 111  
**sgWrite** — Read data from a segment into an actor address space; Write data from an actor address space into a segment 102, 109

## V

**vmFlush** — flush local cache(s) 39, 47, 99, 106, 111